arXiv:1602.04036v1 [math.NA] 12 Feb 2016

A CG-type method in Banach spaces with an application to computerized tomography

Frederik Heber*, Frank Schöpfer†, and Thomas Schuster‡

February 15, 2016

Abstract

Conjugate Gradient (CG) methods are one of the most effective iterative methods to solve linear equations in Hilbert spaces. So far, they have been inherently bound to these spaces since they make use of the inner product structure. In more general Banach spaces one of the most prominent iterative solvers are Landweber-type methods that essentially resemble the Steepest Descent method applied to the normal equation. More advanced are subspace methods that take up the idea of a Krylov-type search space, wherein an optimal solution is sought. However, they do not share the conjugacy property with CG methods. In this article we propose that the Sequential Subspace Optimization (SESOP) method can be considered as an extension of CG methods to Banach spaces. We employ metric projections to orthogonalize the current search direction with respect to the search space from the last iteration. For the ℓ_2-space our method then exactly coincides with the Polak-Ribière type of the CG method when applied to the normal equation. We show that such an orthogonalized search space still leads to weak convergence of the subspace method. Moreover, numerical experiments on a random matrix toy problem and 2D computerized tomography on ℓ_p-spaces show superior convergence properties over all p compared to non-orthogonalized search spaces. This especially holds for ℓ_p-spaces with small p. We see that the closer we are to an ℓ_2-space, the more we recover of the conjugacy property that holds in these spaces, i. e., as expected, the more the convergence behaves independently of the size of the truncated search space.

Keywords: metric projection, Bregman distance, Banach space, sequential subspace optimization methods, conjugate gradient, computerized tomography

AMS: 46B20, 65K10, 65F10, 65R32

*Department of Mathematics, Saarland University, 66123 Saarbrücken, Germany (Email: heber@math.uni-sb.de)

†Department of Mathematics, University of Oldenburg, 26129 Oldenburg, Germany (Email: frank.schoepfer@uni-oldenburg.de)

‡Department of Mathematics, Saarland University, 66123 Saarbrücken, Germany (Email: thomas.schuster@num.uni-sb.de), correspondent author

1 Introduction

We consider two Banach spaces $\mathcal{X}$ and $\mathcal{Y}$ with a continuous linear operator

$$\mathbf{A} : \mathcal{X} \to \mathcal{Y}, \tag{1}$$

where the goal is to (iteratively) solve the inverse problem

$$\mathbf{A}\mathbf{x} = \mathbf{y}. \tag{2}$$

$\mathcal{X}$ is assumed to be smooth and uniformly convex and hence $\mathcal{X}$ is reflexive and has a strictly convex and uniformly smooth dual $\mathcal{X}^*$. $\mathcal{Y}$ can be arbitrary. The problem (2) may be ill-posed and thus not suitable for direct inversion of the operator. Hence, a regularization scheme is required to obtain a stable solution.

In [19] the Landweber method, well-known and thoroughly investigated in Hilbert spaces, see references therein, has been extended to this setting. It is essentially a steepest descent method on the normal equation

$$\mathbf{A}^*\mathbf{A}\mathbf{x} = \mathbf{A}^*\mathbf{y} \tag{3}$$

that results in the minimum-norm solution. However, the method usually suffers from tremendously slow convergence.

Therefore, in [21] the idea of projection onto subspaces has been taken from the family of CG methods. There, multiple simultaneous search directions are admitted to speed up the convergence. These are the so-called subspace methods that differ by the specific choice of search directions. A convex, differentiable line search functional is derived by considering Bregman projections onto hyperplanes that can be efficiently minimized by standard optimization methods, see [15]. This is the Sequential Subspace Optimization (SESOP) method, proposed by [14], where also weak convergence has been shown. Later, by taking specific search directions into account also strong convergence has been proven, see [20, Prop. 1].

However, the set of search directions is still not optimal. In the thesis of [18] the search directions are further modified in the notion of maximizing their pair-wise orthogonality. In this article we give its detailed derivation, which also connects the SESOP method with the Conjugate Gradient (CG) family of methods in Hilbert spaces. Moreover, we propose that the subspace methods also generalize the CG methods from Hilbert spaces to Banach spaces.

In the following we first review some general properties of duality mappings, uniformly smooth Banach spaces, and Bregman distances. Next, we repeat the SESOP method and introduce semi-orthogonalized subspaces. Then, we prove weak convergence of the method with these semi-orthogonalized search space consisting of the last N Landweber descent directions. We conclude with numerical experiments on a random matrix toy problem and 2D computerized tomography to exemplify the significantly improved convergence compared to SESOP [21] with no orthogonalized search spaces.

2 Preliminaries

Throughout the paper let $\mathcal{X}$ and $\mathcal{Y}$ be real Banach spaces with duals $\mathcal{X}^*$ and $\mathcal{Y}^*$. The space $\mathcal{X}$ shall be smooth and uniformly convex with a sequentially

weak-to-weak continuous duality mapping. The space $\mathcal{Y}$ is on the other hand arbitrary. Their norms will be denoted by $\|.\|_{\mathcal{X}}$ and $\|.\|_{\mathcal{X}^*}$, respectively. For $\mathbf{x} \in \mathcal{X}$ and $\mathbf{x}^* \in \mathcal{X}^*$, we write $\langle \mathbf{x}, \mathbf{x}^* \rangle = \langle \mathbf{x}^*, \mathbf{x} \rangle = \mathbf{x}^*(\mathbf{x})$. By $\mathcal{L}(\mathcal{X}, \mathcal{Y})$ we denote the space of all continuous linear operators $\mathbf{A} : \mathcal{X} \to \mathcal{Y}$ and write $\mathbf{A}^*$ for its dual operator $\mathbf{A}^* \in \mathcal{L}(\mathcal{Y}^*, \mathcal{X}^*)$ and $\|\mathbf{A}\| = \|\mathbf{A}^*\|$ for the operator norm of $\mathbf{A}$.[1]

For real numbers a, b, we write

$$a \vee b = \max\{a, b\}, \quad a \wedge b = \min\{a, b\}.$$

Also, let $p, p^*, r, r^* \in (1, \infty)$ be conjugate exponents so that

$$\frac{1}{p} + \frac{1}{p^*} = 1 \quad \text{and} \quad \frac{1}{r} + \frac{1}{r^*} = 1.$$

Let us further note by p and p^* the conjugate exponents satisfying

$$\tfrac{1}{p} + \tfrac{1}{p^*} = 1 \Leftrightarrow 1 - \tfrac{1}{p} = \tfrac{p-1}{p} = \tfrac{1}{p^*} \Leftrightarrow \tfrac{p}{p-1} = p^* \Leftrightarrow \tfrac{p}{p^*} = p - 1 \Leftrightarrow p + p^* = p^* p, \quad (4)$$

where all of these hold also with p and p^* exchanged due to symmetry.

Let us briefly motivate the need for the definitions to follow and also explain the general iterative scheme: We want to find an element $\mathbf{x}$ in a space $\mathcal{X}$ fulfilling (2). In order to assess distances between elements we require a metric or simpler a norm. We will look at a sequence $\{\mathbf{x}_n\}_n$ of iterates $\mathbf{x}_n$ in the space that will converge against the true solution $\widehat{\mathbf{x}}$. We desire this fixed point $\widehat{\mathbf{x}}$ of the iteration to be contained in that space, i. e. it must be closed with respect to the norm. Hence, the space $\mathcal{X}$ we are looking at is a *Banach space*.

Convergence will be proven with respect to the *Bregman distance* and not with respect to the norm, where proof attempts have failed so far. The Bregman distance measures the non-linearity gap between two elements. A further essential ingredient for a proof of convergence is a geometrical inequality that relates the distance between two vectors $\|\mathbf{x} - \mathbf{y}\|$ and the norm of each vector $\|\mathbf{x}\|$ and $\|\mathbf{y}\|$, respectively. These are given by the *Xu-Roach inequalities* in uniformly smooth Banach spaces.

We will provide an update formula that yields the next iterate $\mathbf{x}_{n+1}$ in the sequence $\{\mathbf{x}_n\}_n$ given the current one $\mathbf{x}_n$. This update formula requires an additional search direction used to modify the current iterate.

The residual $(\mathbf{A}\mathbf{x}_n - \mathbf{y})$ is a natural choice for such a search direction as it states how much we are off the true solution $\widehat{\mathbf{x}}$ in the operator's range: $\mathbf{A}\mathbf{x}_n - \mathbf{y} = \mathbf{A}(\mathbf{x}_n - \widehat{\mathbf{x}})$. As we are only given $\mathbf{y}$ in the problem setting, this is all we can do. The residual, however, cannot be used right away as a gradient direction as it resides in $\mathcal{Y}$ and not in $\mathcal{X}$. Now, if the Banach space is not a Hilbert space, then the space and its dual are not isometrically isomorph to one another. In such a case, if the operator $\mathbf{A}$ maps an element $\mathbf{x} \in \mathcal{X}$ to an element $\mathbf{y}$ in a different space $\mathcal{Y}$, then its dual operator $\mathbf{A}^*$ will not map back in the same spaces but between their respective dual spaces, see Figure 1.

To this end, we need operators mapping from a space to its dual and back. Furthermore, these operators should be connected to the norms of the spaces, as the minimum-norm requirement ensures us uniqueness of the solution, see (3). This role is fulfilled by so-called *duality mappings*, see Figure 1 that can be

[1] In a finite-dimensional setting, $\mathbf{A}$ can be represented as a matrix with its adjoint given by the transposed matrix.

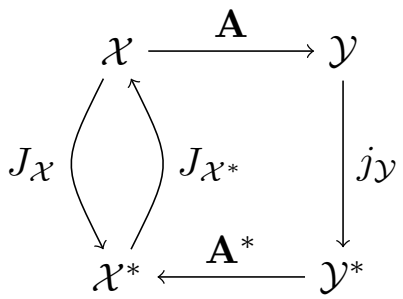

Figure 1: Relations between the Banach spaces and dual spaces involved in the inverse problem $\mathbf{Ax} = \mathbf{y}$.

considered a generalization of a norm's gradient, namely the *subgradient*, to the case when it is not differentiable.

Finally, for a gradient method a step width must be known that scales the current search direction. The Landweber method uses a fixed stepwidth but several dynamic, mostly residual-dependent widths are known as well. Naturally, these are either very expensive to calculate or not optimal. Alternatively, for the subspace methods a line search functional is minimized, where even multiple search directions are admitted. This functional consists of computing a *Bregman projection* onto the intersection of planes defined by the search directions. While multiple search directions improve convergence, they are not related in any special way to one another.

This is then where this article continues: we propose to further use metric projections to semi-*orthogonalize* search directions to enhance the convergence's effectiveness and to keep a memory of outdated search directions. This essentially is the conjugacy property of the CG methods.

2.1 Duality Mappings

We recall the definition of duality mappings and some of their properties, all of which can be found in the comprehensive book [4].

Definition 1 (Duality Mapping)**.** The mapping $J_p : \mathcal{X} \to 2^{\mathcal{X}^*}$ defined by

$$J_p(\mathbf{x}) = \{\mathbf{x}^* \in \mathcal{X}^* \mid \langle \mathbf{x}^*, \mathbf{x} \rangle = \|\mathbf{x}\|^p, \|\mathbf{x}^*\| = \|\mathbf{x}\|^{p-1}\} \tag{5}$$

is the *duality mapping* of $\mathcal{X}$ with gauge function $t \mapsto t^{p-1}$.

We remark that the sets $J_p(\mathbf{x})$ are non-empty by the Hahn-Banach theorem and that they are in general set-valued. We refer to [21] for examples for Hilbert, L_p-spaces, and ℓ_p-spaces. By the theorem of Asplund, see [4, Thm. 4.4], duality mappings are the subdifferential of the norm in the form

$$J_p(\mathbf{x}) = \partial \left(\tfrac{1}{p} \|\mathbf{x}\|^p \right), \tag{6}$$

where the subdifferential is the set of subgradients $\mathbf{x}^*$ of a function f that each fulfil

$$f(\mathbf{y}) - f(\mathbf{x}) \geq \langle \mathbf{x}^*, \mathbf{y} - \mathbf{x} \rangle \quad \text{for all } \mathbf{y} \in \mathcal{X}. \tag{7}$$

That is the theorem of Asplund associates the norm's (sub)gradient in the form (6) with the duality mapping of its space.

Proposition 1 (Single-valued duality mapping, [4, Thm. 3.5, 4.5])**.** $\mathcal{X}$ is smooth if and only if its norm is Gâteaux differentiable on $\mathcal{X} \setminus 0$ iff J_p is single-valued.

2.2 Uniform Smoothness

We will show the convergence of SESOP with orthogonalized search directions in uniformly smooth spaces. The definition of uniform smoothness hinges on the the modulus of smoothness.

Definition 2 (Uniformly smooth, [10, Def. II.1.e.1])**.** Let $\mathcal{X}$ be a Banach space with $\dim \mathcal{X} \geq 2$.

(a) Its *modulus of smoothness* is defined by

$$\rho_{\mathcal{X}}(\tau) = \tfrac{1}{2} \sup_{\|\mathbf{x}\|=1, \|\mathbf{y}\|=1} (\|\mathbf{x} + \tau\mathbf{y}\| + \|\mathbf{x} - \tau\mathbf{y}\| - 2) \quad \tau > 0. \tag{8}$$

(b) $\mathcal{X}$ is said to be uniformly smooth if

$$\lim_{\tau \to \infty} \frac{\rho_{\mathcal{X}}(\tau)}{\tau} = 0. \tag{9}$$

Note that L_p-spaces and ℓ_p-spaces with $p \in (1, \infty)$ are uniformly convex, see [5], and thereby also uniformly smooth, see [6, Lemma 6.7], see also [11, p. 63]. On the other hand, ℓ_1 and l_∞ are not reflexive[2], see [17, Sect. VIII. 5], and hence can be neither uniformly smooth nor uniformly convex.

In order to prove convergence we will heavily rely on the geometrical characteristics of Banach spaces. Essentially, we need to estimate the distance in the norm between two vectors $\|\mathbf{x} - \mathbf{y}\|$ from both norms $\|\mathbf{x}\|, \|\mathbf{y}\|$ alone. These characteristics are captured in the Xu-Roach inequalities, one of which we repeat here for convenience from [18].

Theorem 1 ([22, Theorem 2, Remark 4])**.** If $\mathcal{X}$ is uniformly smooth, then for all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$, we have

$$\|\mathbf{x} - \mathbf{y}\|^p \leq \|\mathbf{x}\|^p - p\langle J_p(\mathbf{x}), \mathbf{y}\rangle + \widetilde{\sigma}_p(\mathbf{x}, \mathbf{y}) \tag{10}$$

with

$$\widetilde{\sigma}_p(\mathbf{x}, \mathbf{y}) = pG_{p^*} \int_0^1 \frac{(\|\mathbf{x} - t\mathbf{y}\| \vee \|\mathbf{x}\|)^p}{t} \rho_{\mathcal{X}}\left(\frac{t\|\mathbf{y}\|}{\|\mathbf{x} - t\mathbf{y}\| \vee \|\mathbf{x}\|}\right) dt \tag{11}$$

where

$$\begin{aligned} G_{p^*} &= 8 \vee 64cK_{p^*}^{-1}, \\ K_{p^*} &= 4(2 + \sqrt{3}) \min\{\tfrac{1}{2}p^*(p^* - 1) \wedge 1, (\tfrac{1}{2}p^* \wedge 1)(p^* - 1), \\ &\quad (p^* - 1)\left[1 - (\sqrt{3} - 1)^p\right], 1 - \left[1 + (2 - \sqrt{3})p\right]^{1-p^*}\}, \end{aligned}$$

and

$$c = \frac{4\tau_0}{\sqrt{1 + \tau_0^2} - 1} \prod_{j=1}^{\infty} \left(1 + \frac{15\tau_0}{2^{j+2}}\right) \quad \text{with } \tau_0 = \frac{\sqrt{339} - 18}{30}. \tag{12}$$

[2] Note that finite-dimensional spaces are always reflexive.

Proof. The proof is given in [22]. Note that we have additionally used $\rho_{\mathcal{X}}(\tau) \geq \sqrt{1+\tau^2} - 1$ for every Banach space $\mathcal{X}$, see [10, p. 243]. □

We also give here a technical lemma on an upper bound on the function $\widetilde{\sigma}_p$ here, which we require later in the convergence proof.

Lemma 1 (Upper bound on $\widetilde{\sigma}$,[18, Proof of Prop. 2.39])**.** Let $\mathcal{X}^*$ be a uniformly smooth Banach space with duality mapping J_{p^*}. If $0 \neq \mathbf{x} \in \mathcal{X}$, $0 \neq \mathbf{A} \in \mathcal{L}(\mathcal{X}, \mathcal{Y})$ and $0 \neq \mathbf{y}^* \in \mathcal{Y}^*$ with an arbitrary Banach space $\mathcal{Y}$ are given and $\mu > 0$ is defined by

$$\mu := \frac{\tau}{\|\mathbf{A}\|} \frac{\|\mathbf{x}\|^{p-1}}{\|\mathbf{y}^*\|} \quad \text{for some } \tau \in (0, 1], \tag{13}$$

then the following estimate is valid:

$$\tfrac{1}{p^*} \widetilde{\sigma}_{p^*}(J_{p^*}(\mathbf{x}), \mu \mathbf{A}^* \mathbf{y}^*) \leq 2^{p^*} G_p \|\mathbf{x}\|^p \rho_{\mathcal{X}^*}(\tau), \tag{14}$$

whereby G_{p^*} is the constant appearing in (11) and $\rho_{\mathcal{X}^*}$ is the modulus of smoothness of $\mathcal{X}^*$, see [4] for definitions and details.

Proof. We look at the definition of (11) for $\widetilde{\sigma}_{p^*}(J_{p^*}(\mathbf{x}), \mathbf{A}^*\mathbf{y}^*)$ and estimate term by term. By (13) we get

$$\|J_{p^*}(\mathbf{x}) - t\mu \mathbf{A}^* \mathbf{y}^*\| \leq \|\mathbf{x}\|^{p-1} + \mu \|\mathbf{A}\| \|\mathbf{y}^*\| \leq 2 \|\mathbf{x}\|^{p-1}$$

and

$$\|J_{p^*}(\mathbf{x}) - t\mu \mathbf{A}^* \mathbf{y}^*\| \vee \|J_{p^*}(\mathbf{x})\| \begin{cases} \leq 2 \|\mathbf{x}\|^{p-1} \\ \geq \|\mathbf{x}\|^{p-1} \end{cases}.$$

As the modulus of smoothness $\rho_{\mathcal{X}^*}$ is non-decreasing, see [11, Prop. 1.e.5], and with (6) and (13), we see that

$$\rho_{\mathcal{X}^*}\left(\frac{t\mu \|\mathbf{A}^*\mathbf{y}^*\|}{\|J_{p^*}(\mathbf{x}) - \mu \mathbf{A}^*\mathbf{y}^*\| \vee \|J_{p^*}(\mathbf{x})\|}\right) \leq \rho_{\mathcal{X}^*}\left(\frac{t\mu \|\mathbf{A}^*\mathbf{y}^*\|}{\|\mathbf{x}\|^{p-1}}\right) \leq \rho_{\mathcal{X}^*}(t\tau).$$

And we finally arrive at the desired estimate,

$$\begin{aligned} \tfrac{1}{p^*} \widetilde{\sigma}_{p^*}(J_{p^*}(\mathbf{x}), \mu \mathbf{A}^* \mathbf{y}^*) &\leq 2^{p^*} G_p \|\mathbf{x}\|^p \int_0^1 \frac{\rho_{\mathcal{X}^*}(t\tau)}{t} dt \\ &= 2^{p^*} G_p \|\mathbf{x}\|^p \int_0^\tau \frac{\rho_{\mathcal{X}^*}(t)}{t} dt \\ &\leq 2^{p^*} G_p \|\mathbf{x}\|^p \rho_{\mathcal{X}^*}(\tau) \end{aligned}$$

as the function $\tau \to \frac{\rho_{\mathcal{X}^*}(\tau)}{\tau}$ is non-decreasing, see [3, Cor. 2.8] □

2.3 Metric and Bregman Projections

We summarize essential properties of metric and Bregman projections. For a projection P onto a closed and convex set $\mathcal{C}$ we have $P^2(\mathbf{x}) = P(\mathbf{x})$ and $P(\mathbf{x}) = \mathbf{x} \Leftrightarrow \mathbf{x} \in \mathcal{C}$.

Definition 3 (Metric Projection)**.** The *metric projection* P of $\mathbf{x} \in \mathcal{X}$ onto $\mathcal{C}$ is the unique element $P_{\mathcal{C}}(\mathbf{x}) \in \mathcal{C}$ such that

$$\|\mathbf{x} - P_{\mathcal{C}}(\mathbf{x})\| = \min_{\mathbf{y}\in\mathcal{C}} \|\mathbf{x} - \mathbf{y}\| . \tag{15}$$

Let us then remind of the Bregman distance in the context of generalized distance functions, see also [2, Sect. 2.1]

Definition 4 (Bregman Distance)**.** For a Gâteaux differentiable convex function $f : \mathcal{X} \to \mathbb{R}$ the function

$$\Delta_f(\mathbf{x}, \mathbf{y}) := f(\mathbf{y}) - f(\mathbf{x}) - \langle f'(\mathbf{x}), \mathbf{y} - \mathbf{x}\rangle, \quad \mathbf{x}, \mathbf{y} \in \mathcal{X} \tag{16}$$

is called the *Bregman distance* of $\mathbf{x}$ to $\mathbf{y}$ with respect to the function f.

Here, we consider Bregman distances of functions $f_p(\mathbf{x}) = \frac{1}{p}\|\mathbf{x}\|^p$ with $f_p' = J_p$, i. e. the duality mapping of $\mathcal{X}$, see (6). A useful identity [21] for the Bregman distance is then

$$\Delta_p(\mathbf{x}, \mathbf{y}) = \frac{1}{q}\|\mathbf{x}\|^p - \langle J_p(\mathbf{x}), \mathbf{y}\rangle + \frac{1}{p}\|\mathbf{y}\|^p . \tag{17}$$

Proposition 2 (Properties of Bregman Distances,[19, Theorem 2.12])**.** For all $\mathbf{x}, \mathbf{y} \in \mathcal{X}$ and sequences $\{\mathbf{x_n}\}_n$ in $\mathcal{X}$ the following holds:

(a) $\Delta_p(\mathbf{x}, \mathbf{y}) \geq 0$ and $\Delta_p(\mathbf{x}, \mathbf{y}) = 0 \Leftrightarrow \mathbf{x} = \mathbf{y}$.

(b) $\lim_{\|x_n\|\to\infty} \Delta_p(\mathbf{x_n}, \mathbf{x}) = \infty$, i. e. the sequence $\{\mathbf{x_n}\}_n$ remains bounded if the sequence $\{\Delta_p(\mathbf{x_n}, \mathbf{x})\}_n$ is bounded.

(c) Δ_p is continuous in both arguments. It is strictly convex and Gâteaux differentiable with respect to the second variable with $\partial_{\mathbf{y}}\Delta_p(\mathbf{x}, \mathbf{y}) = J_p(\mathbf{y}) - J_p(\mathbf{x})$.

It is easy to see that in Hilbert spaces metric distance and Bregman distance coincide.

Finally, equivalent to metric projections in definition 3, Bregman projections minimize the Bregman distance with respect to a given convex, non-empty set.

Definition 5 (Bregman Projection)**.** The *Bregman projection* of $\mathbf{x} \in \mathcal{X}$ onto $\mathcal{C}$ with respect to the function $f_p(\mathbf{x}) = \frac{1}{p}\|\mathbf{x}\|_{\mathcal{X}}^p$ is the unique element $\Pi_{\mathcal{C}}^p(\mathbf{x}) \in \mathcal{C}$ such that

$$\Delta_p\left(\mathbf{x}, \Pi_{\mathcal{C}}^p(\mathbf{x})\right) = \min_{\mathbf{y}\in\mathcal{X}} \Delta_p(\mathbf{x}, \mathbf{y}). \tag{18}$$

Finally, let us look at equalities and differences between these two projections.

Proposition 3 ([21, Prop. 3.6])**.**

(a) The Bregman projection and the metric projection are related via

$$P_{\mathcal{C}}(\mathbf{x}) - \mathbf{x} = \Pi_{\mathcal{C}-\mathbf{x}}^p(0) \quad \text{for all } \mathbf{x} \in \mathcal{X}. \tag{19}$$

Especially we have $P_{\mathcal{C}}(0) = \Pi_{\mathcal{C}}^p(0)$.

(b) The metric projection has the translational property

$$P_{\mathbf{y}+\mathcal{C}}(\mathbf{x}) = \mathbf{y} + P_{\mathcal{C}}(\mathbf{x}-\mathbf{y}) \quad \text{for all } \mathbf{x},\mathbf{y} \in \mathcal{X}. \tag{20}$$

This property indeed distinguishes the metric from the Bregman projection since if we had $\Pi^p_{\mathbf{y}+\mathcal{C}}(x) = \mathbf{y} + \Pi^p_{\mathcal{C}}(\mathbf{x}-\mathbf{y})$ for all $\mathbf{x},\mathbf{y} \in \mathcal{X}$, then this would already imply their equivalence, $\Pi^p_{\mathbf{y}+\mathcal{C}}(\mathbf{x}) = P_{\mathbf{y}+\mathcal{C}}(\mathbf{x})$ for all $\mathbf{x},\mathbf{y} \in \mathcal{X}$.

3 Methods

We now discuss the sequential subspace methods for solving problem (2) given an operator (1) and two Banach spaces $\mathcal{X}$, $\mathcal{Y}$ and their dual spaces $\mathcal{X}^*$, $\mathcal{Y}^*$ with properties as stated in the beginning.

Let us have first a comment on the notation: In the preliminaries section we stated that space $\mathcal{X}$ and its dual space $\mathcal{X}^*$ are not isometrically isomorph. In consequence, we exploit in the following method the relations between spaces as given in Figure 1. To this end, if $\mathbf{v}^*$ is an element in the space $\mathcal{X}^*$ and can be written as $\mathbf{A}^*\mathbf{o}^*$, then we call $\mathbf{o}^*$ in the following the "precursor" of $\mathbf{v}^*$ because of the intimate connection between the spaces $\mathcal{Y}^*$ and $\mathcal{X}^*$ via the adjoint operator $\mathbf{A}^*$.

Of special importance is the optimality condition, see [19, Lemma 2.10], which we repeat here.

Lemma 2 (Optimality condition)**.** Let $\mathcal{X}$ be smooth and uniformly convex and $\mathbf{y} \in \operatorname{ran}(\mathbf{A})$.

(a) There exists the minimum-norm-solution $\widehat{\mathbf{x}}$ of (2) and $J_p(\widehat{\mathbf{x}}) \in \overline{\operatorname{ran}(\mathbf{A}^*)}$.

(b) If $\widehat{\mathbf{x}} \in \mathcal{X}$ is the minimum-norm-solution of (2) and $\widetilde{\mathbf{x}} \in \mathcal{X}$ fulfils $J_p(\widetilde{\mathbf{x}}) \in \overline{\operatorname{ran}(\mathbf{A}^*)}$ and $\widehat{\mathbf{x}} - \widetilde{\mathbf{x}} \in \operatorname{nul}(\mathbf{A})$, then $\widetilde{\mathbf{x}} = \widehat{\mathbf{x}}$.

3.1 Sequential Subspace Optimization

For convenience let us first recall the SESOP method as given in [21] for solving the ill-posed inverse problem $\mathbf{A}\mathbf{x} = \mathbf{y}$ without noise.

Method 1 (SESOP).

(S1) Take $\mathbf{x}_0$ as initial value with $J_p(\mathbf{x}_0) \in \overline{\operatorname{ran}(\mathbf{A}^*)}$, set $n := 0$, $\mathcal{U}_{-1} := \{0\}$ and repeat the following steps:

(S2) If $R_n := \|\mathbf{R}_n\| := \|\mathbf{A}\mathbf{x}_n - \mathbf{y}\| = 0$ then STOP else goto (S3).

(S3) Choose the search space $\mathcal{U}_n = \operatorname{span}\{\mathbf{u}^*_{n,1}, \dots, \mathbf{u}^*_{n,N_n}\} \subset \operatorname{ran}(\mathbf{A}^*)$ with N_n search directions $\mathbf{u}^*_{n,k} \in \mathcal{U}_n$, $k = 1, \dots, N_n$ and with N_n offsets $\alpha_{n,k} := \langle \mathbf{u}^*_{n,k}, \mathbf{z} \rangle$ for any $\mathbf{z} \in \mathcal{M}_{\mathbf{A}\mathbf{x}=\mathbf{y}} := \{\mathbf{x} \in \mathcal{X} : \mathbf{A}\mathbf{x} = \mathbf{y}\}$.

(S4) Compute the new iterate:

$$\mathbf{x}_{n+1} := J_{p^*}\left(J_p(\mathbf{x}_n) - \sum_{k=1}^{N_n} \mu_{n,k} \mathbf{u}^*_{n,k} \right) \tag{21}$$

where $\mu_n = (\mu_{n,1}, \dots, \mu_{n,N_n})$ is the solution of the N_n-dimensional optimization problem

$$\min_{t \in \mathbb{R}^{N_n}} h_n(t)$$

with

$$h_n(t) := \frac{1}{p^*} \left\| J_p(\mathbf{x}_n) - \sum_{k=1}^{N_n} t_k \mathbf{u}_{n,k}^* \right\|^{p^*} + \sum_{k=1}^{N_n} t_k \alpha_{n,k} \tag{22}$$

$$\partial_j h_n(t) = -\left\langle \mathbf{u}_{j,k}^*, J_{p^*}\left(J_p(\mathbf{x}_n) - \sum_{k=1}^{N_n} t_k \mathbf{u}_{n,k}^* \right) \right\rangle + \alpha_{j,k} \quad \forall j = 1, \dots, N_n \tag{23}$$

(S5) Set $n \leftarrow n+1$ and goto (S2).

Note that (22) is strictly convex, hence μ is unique.

Convergence of the method essentially depends on the choice of the search space $\mathcal{U}_n$ and associated offsets α_n per iteration step n, see step (S3). We state a few common choices, taken from [21, 20], using the *Landweber descent direction* $\mathbf{d}_n^* := \mathbf{A}^*\mathbf{R}_n^* = \mathbf{A}^* J_r\big(\mathbf{A}\mathbf{x}_n - \mathbf{y}\big)$ with precursor $\mathbf{R}_n^* = J_r \mathbf{R}_n$. Here J_r denotes a single-valued selection of the set-valued duality mapping of Y.

(a) Expanding: $\mathcal{U}_n^{\text{exp}} = \operatorname{span}\{\mathbf{d}_0^*, \dots, \mathbf{d}_n^*\}$, $\alpha_{n,k}^{\text{exp}} = \langle \mathbf{R}_{n,k}^*, \mathbf{y} \rangle$ with dimension $|\mathcal{U}_n| = n+1$

(b) Truncated: $\mathcal{U}_n^{\text{trunc}} = \operatorname{span}\{\mathbf{d}_{n-N_n+1}^*, \dots, \mathbf{d}_n^*\}$, $\alpha_{n,k}^{\text{trunc}} = \langle \mathbf{R}_{n,k}^*, \mathbf{y} \rangle$ with dimension $|\mathcal{U}_n| = N_n := N \wedge (n+1)$ for some fixed $N \in \mathbb{N}$

(c) Nemirovsky I: $\mathcal{U}_n^{\text{Nem1}} = \operatorname{span}\{\mathbf{d}_n^*, J_{\mathcal{X}}(x_n) - J_{\mathcal{X}}(x_0)\}$, $\alpha_{n,k}^{\text{Nem1}} = \langle \mathbf{v}_{k,l}^*, \mathbf{y} \rangle$ with dimension $|\mathcal{U}_n| = 2$

(d) Nemirovsky II: $\mathcal{U}_n^{\text{Nem2}} = \operatorname{span}\{\mathbf{d}_n^*, J_{\mathcal{X}}(x_n) - J_{\mathcal{X}}(x_{n-1})\}$, $\alpha_{n,k}^{\text{Nem2}} = \langle \mathbf{v}_{k,l}^*, \mathbf{y} \rangle$ with dimension $|\mathcal{U}_n| = 2$

Note that the Nemirovsky directions of cases (c) and (d) provide strong convergence, [20, Prop. 1], are not considered in the scope of this article. Also, for finite-dimensional spaces $\mathcal{X}$ and $\mathcal{Y}$ weak and strong convergence coincide [17, Thm 4.3].

In the cases (a) or (b) the hyperplane offsets can simply be calculated by $\alpha_{n,k} := \langle \mathbf{w}_{n,k}^*, \mathbf{y} \rangle$.

3.2 Generalized CG

Employing multiple search directions greatly improves convergence as indicated by experiments [21, Sect. 5]. However, the search directions are still not related in any particular way among one another. In the following we would like to maximize "distinctiveness" of the search directions $\mathbf{u}_{n,k}^*$ in the truncated search space $\mathcal{U}_n$, i. e. we orthogonalize them with respect to previous ones contained in $\mathcal{U}_n$ and to some extent to older search directions not present in $\mathcal{U}_n$.

Because of the Banach space structure this orthogonality holds only one way: new descent directions are made orthogonal to ones already contained in the search space but not the other way round, i. e. it is not symmetric.

To make such a (semi-)orthogonalized search space distinguishable from $\mathcal{U}_n^{\text{trunc}}$, we denote it as $\mathcal{V}_n^{\text{trunc}}$ in the following. Note that in this section we extend the details connecting SESOP and CG as stated in [18, Sect. 2.6].

Based on this semi-orthogonality, we then want to construct a search space, similar to SESOP's truncated search space $\mathcal{U}_n^{\text{trunc}}$, using orthogonalized directions. To this end, we examine search directions derived in the following way from the Landweber descent direction $\mathbf{d}_n^* = \mathbf{A}^* J_r(\mathbf{A}\mathbf{x}_n - \mathbf{y})$ to obtain the semi-orthogonalized search space $\mathcal{V}_n^{\text{trunc}} := \{\mathbf{v}_{n,k}^*\}_{k=1}^{N_n}$,

$$\mathbf{v}_{n,k}^* := \mathbf{v}_{n-1,k+1}^* \quad \text{for } k = 1, \dots, N_n - 1 \tag{24}$$

$$\mathbf{v}_{n,N_n}^* := \mathbf{d}_n^* - \sum_{k=1}^{N_{n-1}} s_{n,k} \mathbf{v}_{n-1,k}^* \tag{25}$$

$$= \mathbf{d}_n^* - P_{\mathcal{V}_{n-1}^{\text{trunc}}}(\mathbf{d}_n^*)\,. \tag{26}$$

Here, $s_{n,k}$ is an orthogonalization coefficient obtained from metric projection of the Landweber descent direction $\mathbf{d}_n^*$ onto the search space $\mathcal{V}_{n-1} = \{\mathbf{v}_{n-1,k}^*\}_{k=1}^{N_{n-1}}$. For convenience, let us define,

$$s_n = (s_{n,1}, \dots, s_{n,N_n}) := \operatorname{argmin}_{s \in \mathbb{R}^{N_n}} g_n(s) \tag{27}$$

and

$$g_n(s) := \left\| \mathbf{d}_n^* - \sum_{i=1}^{N_{n-1}} s_i \mathbf{v}_{n-1,i}^* \right\|_{\mathcal{X}^*}^{p^*}. \tag{28}$$

In Hilbert spaces this would become the familiar Gram-Schmidt procedure, see also [13, p. 218].

We show that all search directions $\mathbf{v}_{n,j}^*$, $j = 1, \dots, N_n$, are pairwise semi-orthogonal to in a certain sense. This is the Banach space's counterpart of the conjugacy property in a Hilbert space when treating the normal equation, c. f. [15, p. 102].

Corollary 1 (Semi-orthogonal search directions)**.** We have

$$\langle \mathbf{v}_{n,j}^*, J_{p^*}(\mathbf{v}_{n,k}^*) \rangle = 0 \quad \forall 1 \le j < k \le N_n. \tag{29}$$

Proof. This follows directly from the optimality condition on $g_n(s)$ for (27),

$$0 = \partial_j g_n(s_{N_n}) = -\left\langle \mathbf{v}_{n-1,j}^*, J_{p^*}\left(\mathbf{d}_n^* - \sum_{k=1}^{N_{n-1}} s_{n,k} \mathbf{v}_{n-1,k}^* \right) \right\rangle \tag{30}$$

$$= -\langle \mathbf{v}_{n-1,j}^*, J_{p^*}\left(\mathbf{v}_{n,N_n}^*\right) \rangle$$

for all $j = 1, \dots, N_{n-1}$. As this holds for all n and by the successive construction (24) of the search spaces $\mathcal{V}_n^{\text{trunc}}$, we have by induction that all search directions in $\mathcal{V}_n^{\text{trunc}}$ are orthogonal with respect to one another. □

Note that this semi-orthogonality becomes a full orthogonality in the case of a Hilbert space $\mathcal{X}$ with a duality mapping J_p of power type $p = 2$, as J_p can

then be identified with the identity. For simplicity, we speak in the following only of orthogonal search directions.

We would like to stress that the summation in (25) is over all directions in the search space $\mathcal{V}_{n-1}^{\text{trunc}}$, including $\mathbf{v}_{n-1,1}^*$, that is not included in $\mathcal{V}_n^{\text{trunc}}$, see (24). Otherwise, the orthogonalization would not change the search space but only modify its spanning vectors. Therefore, we clearly have the case of $\mathcal{U}_n^{\text{trunc}} \not\subset \mathcal{V}_n^{\text{trunc}}$ for $n \geq 1$ in general. Using the expanding search space $\mathcal{U}_n^{\text{exp}}$ on the other hand, see case (a) in section 3.1, the above orthogonalization of the new search direction would *not* change the iteration, i. e. for every n we have $\mathcal{U}_n^{\text{exp}} = \mathcal{V}_n^{\text{exp}}$.

Example 1. In order to highlight the notational equivalence with CG methods in Hilbert spaces, let us consider briefly just a single search direction, i. e. $\mathcal{V}_n^{\text{trunc}} = \text{span}\{\mathbf{v}_n^*\}$, per iteration step n, we obtain

$$\mathbf{v}_n^* = \mathbf{d}_n^* - s_n \mathbf{v}_{n-1}^*. \tag{31}$$

Here, we have the Landweber descent direction $\mathbf{d}_n^*$ that is modified by the last search direction scaled by the orthogonalization coefficient, i. e. $\mathbf{g_n} = \mathbf{d_n} - s_n \mathbf{g_{n-1}}$ where $\mathbf{g_n}$ is the current search direction and $\mathbf{d_n}$ is the current gradient direction in the usual notation of the CG methods in Hilbert spaces, c. f. [15, Chap. 5] and also [23].

With each search direction a hyperplane offset is required that relates this hyperplane to the solution manifold $\mathcal{M}_{\mathbf{Ax=y}}$, c. f. step (S3) in method 1. Naturally, the last offset changes this step when we orthogonalize the descent direction $\mathbf{d}_n^*$. Hence, instead of α_{n,N_n} associated to $\mathbf{d}_n^*$, we require a β_{n,N_n} associated to $\mathbf{v}_n^*$.

Calculating the new hyperplane offsets $\beta_{n,k}$ to each orthogonalized search direction $\mathbf{v}_{n,k}^*$ in $\mathcal{V}_n^{\text{trunc}}$ is then,

$$\beta_{n,N_n} := \langle \mathbf{o}_{n,N_n}^*, \mathbf{y} \rangle = \langle \mathbf{R}_n^* - \sum_{k=1}^{N_{n-1}} s_{n,k} \mathbf{o}_{n-1,k}^*, \mathbf{y} \rangle = \alpha_{n,N_n} - \sum_{k=1}^{N_{n-1}} s_{n,k} \beta_{n-1,k}, \tag{32}$$

where we used

$$\begin{aligned} \mathbf{v}_{n,N_n}^* &= \mathbf{d}_n^* - \sum_{k=1}^{N_{n-1}} s_{n,k} \mathbf{v}_{n-1,k}^* = \mathbf{d}_n^* - \sum_{k=1}^{N_{n-1}} s_{n,k} \mathbf{A}^* \mathbf{o}_{n-1,k}^* \\ &= \mathbf{A}^* \left(\mathbf{R}_n^* - \sum_{k=1}^{N_{n-1}} s_{n,k} \mathbf{o}_{n-1,k}^* \right) =: \mathbf{A}^* \mathbf{o}_{n,N_n}^*, \end{aligned} \tag{33}$$

with the precursor $\mathbf{R}_n^* = J_r \mathbf{R}_n$ of the Landweber descent direction $\mathbf{R}_n^*$, $\mathbf{d}_n^* = \mathbf{A}^* \mathbf{R}_n^*$.

These orthogonalization coefficients $s_{n,k}$ are calculated by minimizing (28) with the derivative (30) using standard techniques. Note that in the proof of Lemma 4 we give a good starting value for this line search problem. Once we know the coefficients $s_{n,k}$ in (25), we can easily evaluate (32), knowing all other offsets $\beta_{n-1,k}$, $k = 1, \ldots, N_{n-1}$, from previous iterations,

3.3 Proof of Convergence

Now we would like to show that, with the semi-orthogonalized truncated search space $\mathcal{V}_n^{\text{trunc}}$, method 1 still converges weakly to a solution of $\mathbf{Ax} = \mathbf{y}$.

We will assume that $R_n \neq 0$ without loss of generality in the following theorem and in some of the corollaries and lemmata in support of this theorem. This is a valid assumption because if we get $R_n = 0$ at some step n, then it holds that $\mathbf{Ax}_n = \mathbf{y}$ and we are done. As a direct consequence of this assumption, it follows that the Landweber descent direction is always non-zero.

Corollary 2 (Non-zero descent direction)**.** If $R_n \neq 0$, we have $\mathbf{d}_n^* \neq 0$.

Proof. Assume the contrary, $\mathbf{d}_n^* = 0$, and let be given a $\mathbf{z} \in \mathcal{M}_{\mathbf{Ax}=\mathbf{y}}$, then we get

$$0 = \langle \mathbf{d}_n^*, \mathbf{x}_n - \mathbf{z} \rangle = \langle \mathbf{R}_n^*, \mathbf{Ax}_n - \mathbf{Az} \rangle = \|\mathbf{R}_n\|^r = R_n^r,$$

which is a contradiction.

Note that by the very same argument we also have $\mathbf{R}_n^* \neq 0$. □

Then, we need to prove something similar to the "expanding subspace" property of the CG methods, c. f. [15, Theorem 5.2].

Corollary 3 (Truncated Subspace Minimization)**.** At iteration step n let be given an orthogonalized search space $\mathcal{V}_n^{\text{trunc}} := \text{span}\{\mathbf{v}_{n,1}^*, \dots, \mathbf{v}_{n,N_n}^*\}$ with precursors $\{\mathbf{o}_{n,i}^*\}_{i=1}^{N_n}$.

(a) We have that all old search directions still contained in $\mathcal{V}_n^{\text{trunc}}$ are orthogonal with respect to the dual pairing with current and old residuals,

$$\begin{aligned} &\langle \mathbf{o}_{n-1,1}^*, \mathbf{R}_{n-N_n} \rangle = 0 \\ &\dots \\ &\langle \mathbf{o}_{n-1,1}^*, \mathbf{R}_n \rangle = 0, \dots, \langle \mathbf{o}_{n-1,N_{n-1}}^*, \mathbf{R}_n \rangle = 0. \end{aligned} \tag{34}$$

Note that this extends to $\{\mathbf{o}_{n,1}^*, \dots, \mathbf{o}_{n,N_n-1}^*\}$ by construction of $\mathcal{V}_n^{\text{trunc}}$.

(b) If we have $R_n \neq 0$, then it also holds that $\mathbf{v}_{n,N_n}^* \neq 0$.

(c) The set of vectors $\{\mathbf{o}_{n,1}^*, \dots, \mathbf{o}_{n,N_n}^*\}$ and $\{\mathbf{v}_{n,1}^*, \dots, \mathbf{v}_{n,N_n}^*\}$ are each linearly independent.

Proof. We first prove (a). Assume $R_n \neq 0$, i. e. $\mathbf{R}_n \neq 0$. Let us inspect the optimality condition of the step width functional $h_{n-1}(\mu)$, see (22), at $n-1$ for $j = 1, \dots, N_{n-1}$,

$$\begin{aligned} 0 &= \partial_j h_{n-1}(\mu_{n-1}) \\ &= -\langle \mathbf{v}_{n-1,j}^*, J_{p^*} \left(J_p(\mathbf{x}_{n-1}) - \sum_{k=1}^{N_{n-1}} \mu_{n-1,k} \mathbf{v}_{n-1,k}^* \right) \rangle + \beta_{n-1,j} \\ &= -\langle \mathbf{o}_{n-1,j}^*, \mathbf{Ax}_n \rangle + \langle \mathbf{o}_{n-1,j}^*, \mathbf{y} \rangle \\ &= -\langle \mathbf{o}_{n-1,j}^*, \mathbf{R}_n \rangle, \end{aligned}$$

where we needed (21) and (32). The statement then follows by (24) and stepping back until $n - N_n$, knowing that

$$\mathcal{V}_n^{\text{trunc}} = \{\mathbf{v}^*_{n-N_n-1,N_{n-N_n}-1}, \dots, \mathbf{v}^*_{n,N_n}\}.$$

Then, we can continue with (b). By Corollary 2 we have $\mathbf{d}_n^* \neq 0$ and thereby with (25) we have to show that $\mathbf{v}^*_{n,N_n}$ is not contained in $\mathcal{V}_{n-1}$. To this end, let be $\mathbf{z} \in \mathcal{M}_{\mathbf{Ax}=\mathbf{y}}$ and thus $\mathbf{x}_n - \mathbf{z} \neq 0$. Furthermore, let be given $\lambda_1, \dots, \lambda_{N_{n-1}}, \sigma \in \mathbb{R}$ with

$$\sum_{k=1}^{N_{n-1}} \lambda_k \mathbf{v}^*_{n-1,k} + \sigma \mathbf{d}_n^* = 0.$$

Then by using (a),

$$\begin{aligned} 0 &= \sum_{k=1}^{N_{n-1}} \lambda_k \langle \mathbf{v}^*_{n-1,k}, \mathbf{x}_n - \mathbf{z}\rangle + \sigma \langle \mathbf{d}_n^*, \mathbf{x}_n - \mathbf{z}\rangle \\ &= \sum_{k=1}^{N_{n-1}} \lambda_k \langle \mathbf{o}^*_{n-1,k}, \mathbf{R}_n\rangle + \sigma \langle J_r(\mathbf{R}_n), \mathbf{R}_n\rangle \\ &= \sigma R_n^r, \end{aligned}$$

we get $\sigma = 0$ by contradiction.

Next, for (c) it suffices to show that $\{\mathbf{v}^*_{n,1}, \dots, \mathbf{v}^*_{n,N_n}\}$ are linearly independent. Assume again $R_n \neq 0$ and let similarly be given $\mathbf{z} \in \mathcal{M}_{\mathbf{Ax}=\mathbf{y}}$ and $\lambda_1, \dots, \lambda_{N_n} \in \mathbb{R}$ with

$$\sum_{k=1}^{N_n} \lambda_k \mathbf{v}^*_{n,k} = 0.$$

Using (25) and (a) on $\mathbf{v}^*_{n,N_n}$ and looking at

$$\begin{aligned} 0 = \sum_{k=1}^{N_n} \lambda_k \langle \mathbf{v}^*_{n,k}, \mathbf{x}_n - \mathbf{z}\rangle &= \sum_{k=1}^{N_n-1} \lambda_k \langle \mathbf{o}^*_{n,k}, \mathbf{R}_n\rangle + \lambda_{N_n} \langle \mathbf{o}^*_{n,N_n}, \mathbf{R}_n\rangle \\ = \lambda_{N_n} \langle \mathbf{o}^*_{n,N_n}, \mathbf{R}_n\rangle = \lambda_{N_n} \langle \mathbf{R}_n^*, \mathbf{R}_n\rangle &- \sum_{k=1}^{N_n} s_{n,k} \langle \mathbf{o}^*_{n-1,k}, \mathbf{R}_n\rangle \\ = \lambda_{N_n} R_n^r, & \end{aligned}$$

we get $\lambda_{N_n} = 0$. We continue with

$$0 = \sum_{k=0}^{N_n-1} \lambda_k \langle \mathbf{v}^*_{n,k}, \mathbf{x}_{n-1} - \mathbf{z}\rangle = \dots = \lambda_{N_n-1} R^r_{n-1},$$

and by induction we get $\lambda_k = 0$ for all $k = 1, \dots, N_n$. Hence, also the search directions $\mathbf{v}^*_{n,k}$ are linearly independent.

□

Last but not least, we show that our solution manifold is contained in the intersection of hyperplanes and that iterates and search directions obtained via

the update formulas (21) and (25) derive from Bregman projections onto the intersection of hyperplanes and search space, respectively. While most of this is not needed in the convergence proof, it is very illustrative for the general procedure.

Lemma 3 (Intersection of hyperplanes)**.**

(a) For the solution manifold $\mathcal{M}_{\mathbf{Ax}=\mathbf{y}}$ it holds

$$\mathcal{M}_{\mathbf{Ax}=\mathbf{y}} \subset \mathcal{H}_n := \bigcap_{k=1}^{N_n} \mathcal{H}(\mathbf{v}^*_{n,k}, \beta_{n,k}) \tag{35}$$

with the intersection of the hyperplanes $\mathcal{H}_n$ of all search directions in $\mathcal{V}_n^{\text{trunc}}$.

(b) Also, we have that $J_p(\mathbf{x}_n) - J_p(\mathbf{x}_0) \in \bigcup_n \mathcal{V}_n \subset \overline{\text{ran}\,(\mathbf{A}^*)}$ for all n.

(c) Furthermore, the next iterate $\mathbf{x}_{n+1}$ is the Bregman projection of the current iterate $\mathbf{x}_n$ onto the intersection,

$$\mathbf{x}_{n+1} = \Pi^p_{\mathcal{H}_n}(\mathbf{x}_n) \tag{36}$$

and also,

$$J_{p^*}(\mathbf{x}_{n+1}) = \Pi^p_{J_p(\mathbf{x}_n)+\mathcal{V}_n^{\text{trunc}}}\left(J_p(\mathbf{z})\right) \quad \forall \mathbf{z} \in \mathcal{M}_{\mathbf{Ax}=\mathbf{y}}. \tag{37}$$

(d) We have the search direction $\mathbf{v}^*_{n,N_n}$ as the Bregman projection of the Landweber descent direction $\mathbf{d}^*_n$,

$$J_{p^*}(\mathbf{v}^*_{n,N_n}) = \Pi^p_{(\mathcal{V}_{n-1}^{\text{trunc}})^\perp}\left(J_{p^*}(\mathbf{d}^*_n)\right), \tag{38}$$

where $\mathcal{V}^\perp$ designates the annihilator of the space $\mathcal{V}$.

Proof. For part (a) for any $\mathbf{z} \in \mathcal{M}_{\mathbf{Ax}=\mathbf{y}}$ we have to show $\langle \mathbf{v}^*_{n,k}, \mathbf{z}\rangle = \beta_{n,k}$, which follows directly from the definition of the offsets, (32). Part (b) follows from

$$\begin{aligned} J_p(\mathbf{x}_n) - J_p(\mathbf{x}_0) &= J_p(\mathbf{x}_n) - J_p(\mathbf{x}_{n-1}) + J_p(\mathbf{x}_{n-1}) - \ldots + J_p(\mathbf{x}_1) - J_p(\mathbf{x}_0), \\ &= \sum_{j=1}^{N_n} \mu_n \mathbf{v}^*_{n,j} + \ldots + \sum_{j=1}^{N_1} \mu_1 \mathbf{v}^*_{1,j}, \end{aligned}$$

c. f. (21),

For (d) we only need to use the definition (25) of $\mathbf{v}^*_{n,N_n}$, relations between metric and Bregman projections, see Prop. 3 (a) and [21, Prop. 3.6 d)], and equivalencies for Bregman projections, see [21, Prop. 3.7 b)],

$$\begin{aligned} J_{p^*}(\mathbf{v}^*_{n,N_n}) &= J_{p^*}\left(\mathbf{d}^*_n - P_{\mathcal{V}_{n-1}^{\text{trunc}}}(\mathbf{d}^*_n)\right) = J_{p^*}\left(-\Pi_{\mathcal{V}_{n-1}^{\text{trunc}}-\mathbf{d}^*_n}(0)\right) \\ &= J_{p^*}\left(\Pi_{\mathbf{d}^*_n-\mathcal{V}_{n-1}^{\text{trunc}}}(0)\right) = J_{p^*}\left(\Pi_{\mathbf{d}^*_n+\mathcal{V}_{n-1}^{\text{trunc}}}(0)\right) \\ &= \Pi_{(\mathcal{V}_{n-1}^{\text{trunc}})^\perp}\left(J_{p^*}(\mathbf{d}^*_n)\right). \end{aligned}$$

Part (c) follows as in the proof of [21, Prop. 4.1]. □

Before we then may prove weak convergence, we need to show that $\mathbf{v}^*_{n,N_n}$ is still a descent direction, where we use the same geometrical arguments as in the generalized Landweber convergence proof, see [19], on $\mathbf{d}^*_n$ being a descent direction. This is not a straight-forward consequence of this as we have $\mathcal{U}_n \not\subset \mathcal{V}_n^{\text{trunc}}$, i. e. $\mathbf{d}^*_n$ is not generally contained in $\mathcal{V}_n^{\text{trunc}}$, c. f. (26).

Lemma 4 (Descent direction property)**.** Any $\mathbf{v}^*_{n,N_n}$ resulting from (25) is always a descent direction, i. e. there is a $\mu_n \in \mathbb{R}^{N_n}$ and $S_n > 0$ with

$$h_n(\mu_n) \leq h_n(0) - S_n, \tag{39}$$

which for any $\mathbf{z} \in \mathcal{M}_{\mathbf{Ax}=\mathbf{y}}$ can also be written as

$$\Delta_{\mathcal{X}}(\mathbf{x}_{n+1}, \mathbf{z}) \leq \Delta_{\mathcal{X}}(\mathbf{x}_n, \mathbf{z}) - S_n. \tag{40}$$

Proof. We assume $R_n \neq 0$ and $\mathbf{x}_n \neq 0$. We set

$$\widetilde{\mu}_n := (0, \dots, 0, \nu_n) \quad \text{with } \nu_n := \frac{\tau_n \|\mathbf{x}_n\|^{p-1}_{\mathcal{X}}}{\left\|\mathbf{v}^*_{n,N_n}\right\|_{\mathcal{X}^*}} \tag{41}$$

with $\tau_n \in (0, 1]$ chosen as to fulfil with $\gamma \in (0, 1)$, c. f. Theorem 1,

$$\frac{\rho_{\mathcal{X}^*}(\tau_n)}{\tau_n} = \rho_{\mathcal{X}^*}(1) \wedge \left(\frac{\gamma}{2^{p^*} G_p} \frac{R^r_n}{\|\mathbf{x}_n\|_{\mathcal{X}} \|\mathbf{v}^*_n\|_{\mathcal{X}^*}} \right).$$

Let $\mu_n = \operatorname{argmin}_{t \in \mathbb{R}^n} h_n(t)$, then we estimate with the Xu-Roach inequality (10) and using $\widetilde{\mu}_n$,

$$\begin{aligned}
h_n(\mu_n) &\leq h_n(\widetilde{\mu}_n) \\
&\leq \tfrac{1}{p^*} \|\mathbf{x}_n\|^p_{\mathcal{X}} - \nu_n \langle \mathbf{v}^*_{n,N_n}, \mathbf{x}_n \rangle + \frac{1}{p^*} \widetilde{\sigma} \left(J(\mathbf{x}_n), \nu_n \mathbf{v}^*_{n,N_n} \right) + \langle \mathbf{o}^*_{n,N_n}, \mathbf{y} \rangle \\
&= \tfrac{1}{p^*} \|\mathbf{x}_n\|^p_{\mathcal{X}} - \nu_n \langle \mathbf{R}^*_n - \sum_{k=1}^{N_{n-1}} s_{n,k} \mathbf{o}^*_{n-1,k}, \mathbf{R}_n \rangle + \frac{1}{p^*} \widetilde{\sigma} \left(J(\mathbf{x}_n), \nu_n \mathbf{v}^*_{n,N_n} \right) \\
&= \tfrac{1}{p^*} \|\mathbf{x}_n\|^p_{\mathcal{X}} - \nu_n R^r_n + \frac{1}{p^*} \widetilde{\sigma} \left(J(\mathbf{x}_n), \nu_n \mathbf{v}^*_{n,N_n} \right) \\
&= \tfrac{1}{p^*} \|\mathbf{x}_n\|^p_{\mathcal{X}} - \tau_n \|\mathbf{x}_n\|^{p-1} \frac{R^r_n}{\left\|\mathbf{v}^*_{n,N_n}\right\|_{\mathcal{X}^*}} + \frac{1}{p^*} \widetilde{\sigma} \left(J(\mathbf{x}_n), \nu_n \mathbf{v}^*_{n,N_n} \right),
\end{aligned}$$

where we have used the orthogonality stated in Corollary 3 (a). As the metric projection is non-expanding, c. f. (26), we have

$$\left\|\mathbf{v}^*_{n,N_n}\right\|_{\mathcal{X}^*} \leq \|\mathbf{A}^* J_r(\mathbf{A}\mathbf{x}_n - \mathbf{y})\| \leq \|\mathbf{A}\| R^{r-1}_n.$$

Lemma 1 allows us to bound the last summand using the requirement on τ_n,

$$\begin{aligned}
\tfrac{1}{p^*} \widetilde{\sigma}_{p^*} \left(J(\mathbf{x}_n), \nu_n \mathbf{v}^*_{n,N_n} \right) &\leq 2^{p^*} G_p \|\mathbf{x}_n\|^p_{\mathcal{X}} \rho_{\mathcal{X}^*}(\tau_n) \\
&\leq \tau_n 2^{p^*} G_p \|\mathbf{x}_n\|^p_{\mathcal{X}} \frac{\rho_{\mathcal{X}^*}(\tau_n)}{\tau_n} \\
&\leq \tau_n 2^{p^*} G_p \|\mathbf{x}_n\|^p_{\mathcal{X}} \frac{\gamma}{2^{p^*} G_p} \frac{R^r_n}{\|\mathbf{x}_n\|_{\mathcal{X}} \left\|\mathbf{v}^*_{n,N_n}\right\|_{\mathcal{X}^*}} \\
&= \gamma \tau_n \|\mathbf{x}_n\|^{p-1}_{\mathcal{X}} \frac{R^r_n}{\left\|\mathbf{v}^*_{n,N_n}\right\|_{\mathcal{X}^*}}.
\end{aligned}$$

All together we obtain

$$h_n(\mu_n) \leq \tfrac{1}{p^*} \|\mathbf{x}_n\|_{\mathcal{X}}^p - (1-\gamma)\tau_n \|\mathbf{x}_n\|_{\mathcal{X}}^{p-1} \frac{R_n^r}{\left\|\mathbf{v}_{n,N_n}^*\right\|}$$
$$\leq \tfrac{1}{p^*} \|\mathbf{x}_n\|_{\mathcal{X}}^p - \frac{(1-\gamma)}{\|\mathbf{A}\|}\tau_n \|\mathbf{x}_n\|_{\mathcal{X}}^{p-1} R_n. \tag{42}$$

We are done as all factors in $S_n := \frac{(1-\gamma)}{\|\mathbf{A}\|}\tau_n \|\mathbf{x}_n\|_{\mathcal{X}}^{p-1} R_n$ are positive and knowing $h_n(0) = \frac{1}{p^*} \|\mathbf{x}_n\|_{\mathcal{X}}^p$.

Lastly, let us show how $h_n(\mu)$ and $\Delta_p(\mathbf{x}_{n+1}(\mu), \mathbf{z})$ are related using (17) and (21). For any $\mathbf{z} \in \mathcal{M}_{\mathbf{Ax}=\mathbf{y}}$ we have

$$\begin{aligned}\Delta_p(\mathbf{x}_{n+1}(\mu), \mathbf{z}) &= \tfrac{1}{p^*} \|\mathbf{x}_{n+1}(\mu)\|_{\mathcal{X}}^p - \langle J_p(\mathbf{x}_n) - \sum_{k=1}^{N_n} \mu_k \mathbf{v}_{n,k}^*, \mathbf{z}\rangle + \tfrac{1}{p} \|\mathbf{z}\|_{\mathcal{X}}^p \\ &= \tfrac{1}{p^*} \|\mathbf{x}_{n+1}(\mu)\|_{\mathcal{X}}^p - \langle J_p(\mathbf{x}_n), \mathbf{z}\rangle + \sum_{k=1}^{N_n} \mu_k \beta_{n,k} + \tfrac{1}{p} \|\mathbf{z}\|_{\mathcal{X}}^p \\ &= h_n(\mu) - \langle J_p(\mathbf{x}_n), \mathbf{z}\rangle + \tfrac{1}{p} \|\mathbf{z}\|_{\mathcal{X}}^p .\end{aligned}$$

As the last two terms are constant with respect to μ, they cancel out when considering the difference $h_n(0) - h_n(\mu)$. Hence, (40) holds. □

Finally, we arrive at the proof of convergence.

Theorem 2 (Weak convergence, truncated, semi-orthogonalized search space)**.** Given a uniformly convex and smooth Banach space $\mathcal{X}$ with sequentially weak-to-weak continuous duality mapping and an arbitrary Banach space $\mathcal{Y}$, then with $1 \leq N_n \leq n$ and search space $\mathcal{V}_n^{\text{trunc}}$, given by (24) and (25), method 1 either stops after a finite number $n \in \mathbb{N}$ of iterations (in case $R_n = 0$) with $\mathbf{x}_n$ being the Bregman projection $\widehat{\mathbf{x}} = \Pi_{\mathcal{M}_{\mathbf{Ax}=\mathbf{y}}}(\mathbf{x}_0)$ of $\mathbf{x}_0$ onto the solution manifold $\mathcal{M}_{\mathbf{Ax}=\mathbf{y}}$ or the sequence of the iterates $\{\mathbf{x}_n\}_n$ converges weakly to $\widehat{\mathbf{x}}$.

Proof. In case $R_{n^*} = 0$ for some n^*, then we have $\mathbf{x}_{n^*} \in \mathcal{M}_{\mathbf{Ax}=\mathbf{y}}$ and we are done by [21, Proposition 3.7 b)] together with the optimality condition in Lemma 2 (b).

Let us then assume $R_n \neq 0$ for all n. Lemma 4 ensures that $\{\Delta_p(\mathbf{x}_n, \mathbf{z})\}_n$ for all $\mathbf{z} \in \mathcal{M}_{\mathbf{Ax}=\mathbf{y}}$ is strictly decreasing. Therefore, $\{\Delta_p(\mathbf{x}_n, \mathbf{z})\}_n$ is bounded from above by $\{\Delta_p(\mathbf{x}_0, \mathbf{z})\}_n$.

Then, Proposition 2 (b) assures that $\{\mathbf{x}_n\}_n$ is bounded. As we require $\mathcal{X}$ to be uniformly convex and hence reflexive by the Milman-Pettis theorem, [4, Sect. II.2, Thm. 2.9], every subsequence of $\{\mathbf{x}_n\}_n$ has in turn a subsequence $\{\mathbf{x}_{n_k}\}_k$ that converges weakly to some $\widetilde{\mathbf{x}} \in \mathcal{X}$, see [17, Chap. 8,Thm 4.2].

The proof of $\{R_{n_k}\}_k$ being a null sequence follows in exactly the same way as in [19, p. 320]. Then, we even have $\widetilde{\mathbf{x}} \in \mathcal{M}_{\mathbf{Ax}=\mathbf{y}}$ as $R_{n_k} = \|\mathbf{A}\mathbf{x}_{n_k} - \mathbf{y}\|_{\mathcal{Y}} \to 0$ with $k \to \infty$, i. e. $\mathbf{A}\widetilde{\mathbf{x}} = \mathbf{y}$.

As $\overline{\operatorname{ran}(\mathbf{A}^*)}$ is convex and norm-closed, it is also weakly closed, see [7, Chap. 5, Thm. 3.13]. This together with Lemma 3 (b) implies $J_p(\widetilde{\mathbf{x}}) - J_p(\mathbf{x}_0) \in \overline{\operatorname{ran}(\mathbf{A}^*)}$ given J_p is sequentially weak-to-weak continuous, see the note below. By the optimality condition Lemma 2 (b) with the requirement $J_p(\mathbf{x}_0) \in$

$\overline{\operatorname{ran}(\mathbf{A}^*)}$ and with $\mathbf{z} + \operatorname{nul}(\mathbf{A}) = \mathcal{M}_{\mathbf{Ax}=\mathbf{y}}$ for all $\mathbf{z} \in \mathcal{M}_{\mathbf{Ax}=\mathbf{y}}$, we then conclude $\widetilde{\mathbf{x}} = \Pi_{\mathcal{M}_{\mathbf{Ax}=\mathbf{y}}}(\mathbf{x}_0)$.

As we have shown that for every subsequence of $\{\mathbf{x}_n\}_n$ there is a subsequence that in turn converges weakly to the same limit $\Pi_{\mathcal{M}_{\mathbf{Ax}=\mathbf{y}}}(\mathbf{x}_0)$, then this is valid for the sequence $\{\mathbf{x}_n\}_n$, too, see [24, Sect. 10.5]. □

Note that the duality mappings of ℓ_p-spaces, where $1 < p < \infty$, are sequentially weak-to-weak-continuous, see [21, Remark 4.3].

4 Experiments

In the following we will perform numerical experiments that compare method 1 with search spaces $\mathcal{U}_n^{\text{trunc}}$ and $\mathcal{V}_n^{\text{trunc}}$. The first part is identical to the experiments in [21, Sect. 5], namely solving $\mathbf{Ax} = \mathbf{y}$ for various ℓ_p spaces and uniformly distributed random matrices $\mathbf{A}$ and right-hand sides $\mathbf{y}$. In the second part we solve inverse problems in 2D computerized tomography with the Radon transform as linear operator.

All experiments habe been performed on a single core of an Intel Xeon E3-1270 cpu with 3.50GHz. Note that SESOP and the above described orthogonalization procedure have been implemented in the C++ library *BASSO* (*BA*nach *S*equential *S*ubspace *O*ptimizer), based on the Eigen3 library [8] for the linear algebra routines. It is available on request from the authors.

4.1 Toy Problem

We first look at the inverse toy problem of a random matrix and a random right-hand side to be formally inverted as is done frequently in the literature, but with well-known short-comings, see [12].

Most importantly in this experiment, we want to check the case of $p = 2$ for a single and for multiple search directions. With SESOP there is still a significant decrease in required iterations from single to multiple search direction for the ℓ_2 space. With CG no such difference arises due to the conjugacy property. Furthermore, to assess a possible speed-up of the orthogonalized search directions, we look at various ℓ_p spaces and norms.

4.1.1 Procedure

To this end, we create a uniformly distributed random matrix $\mathbf{A} \in [-1, 1]^{l \times m}$ with $l = 1000$ and $m = 5000$, representing a discretized version of some random operator.

Next, we want to create a "solution" $\widehat{\mathbf{x}} \in \mathbb{R}^m$ to a random right-hand side $y \in \mathbb{R}^l$ in the sense that this solution should be a minimum-norm solution in an ℓ_p-space with $p \in \{1.1, 1.2, 1.5, 2, 3, 6, 10\}$. Therefore, we create a random right-hand side precursor $y^* \in [-1, 1]^l$ and calculate the minimum-norm solution $x^\dagger$ as follows

$$\widehat{\mathbf{x}} := \frac{J_{p^*}^*(\mathbf{A}^* y^*)}{\left\| J_{p^*}^*(\mathbf{A}^* y^*) \right\|_p} \tag{43}$$

And then we finally obtain the right-hand side as $\mathbb{R}^l \ni y = \mathbf{A}\widehat{\mathbf{x}}$.

Table 1: Comparison of SESOP implementations from [21] and this work (without orthogonalization): Average iteration counts n and standard deviations σ_n for various p and N values and matrix dimension $m = 5000$.

(a) from [21]

p	N	n
1.2	2	435
1.2	4	211
1.2	6	137
1.5	2	22
1.5	4	15
1.5	6	14
6	2	102
6	4	79
6	6	49
10	2	297
10	4	183
10	6	131

(b) This work

p	N	n	σ_n
1.2	2	402.9	80.75
1.2	4	199.1	35.72
1.2	6	133.2	19.59
1.5	2	21.7	0.46
1.5	4	14.9	0.3
1.5	6	14	0
6	2	691.7	744.38
6	4	200.1	90.35
6	6	99.2	31.99
10	2	1,413.8	1,399.17
10	4	396.6	195.55
10	6	188.3	64.18

We use 10 different seeds $s \in \{420, \dots, 429\}$ for the random number generator and calculate average iteration counts n and standard deviations σ_n over these 10 runs with otherwise identical parameters. The iteration is stopped at either $n > 20.000$ or if the relative residual $\frac{\|\mathbf{A}\mathbf{x}_n - \mathbf{y}\|_{\mathcal{Y}}}{\|\mathbf{y}\|_{\mathcal{Y}}}$ drops below 10^{-4}.

Note that we always set $\mathcal{Y} = \ell_2(\mathbb{R}^l)$.

4.1.2 SESOP

First, we reproduce the results from [21, Sect. 5] to elucidate any possible differences that arise from different implementations, see Table 1. There, the matrix dimension is 1000×5000, the ℓ_p-norms of $\mathcal{X}$ are given by $p \in \{1.2, 1.5, 6, 10\}$, the power of the gauge function of the duality mapping J_p is given by $\begin{cases} p, p \geq 2 \\ 2, \text{else} \end{cases}$ and we use search direction numbers $N \in \{2, 4, 6\}$.

Taking into account that in Table 1(a) only a single run is given and comparing this to our averages and standard deviations, the iteration counts are in very good agreement up until $p = 2$. For larger p the discrepancy is quite large. However, this is also true for the standard deviations.

Hence, overall we do not see any discrepancy resulting from the different implementations, i. e. we have a solid base for comparing the results with the original MatLab implementation of [21].

4.1.3 Orthogonalized Search Directions

Next, we look at the change in iteration counts between the search space $\mathcal{U}_n^{\text{trunc}}$ used in SESOP and the search spaces $\mathcal{V}_n^{\text{trunc}}$ using metric projections proposed in this article. We use various numbers of search directions $N \in \{1, 2, 4, 6\}$ and

Table 2: Averaged iteration counts n and standard deviations σ_n for various p and N values and $m = 5000$ for solving with SESOP for the minimum-norm solution of $\mathbf{Ax} = \mathbf{y}$: *unorth*ogonalized refers to $\mathcal{U}^{\text{trunc}}$ and *metric* projection to $\mathcal{V}^{\text{trunc}}$.

p	N	n^{unorth}	σ_n^{unorth}	n^{metric}	σ_n^{metric}
1.1	1	18,880.6	1,533.87	2,196.4	828.39
1.1	2	13,739.7	3,078.62	1,243	252.83
1.1	4	7,167.5	1,004.05	815.6	67
1.1	6	4,187.3	509.16	654.4	65.69
1.2	1	627.8	132.57	77.7	12.4
1.2	2	402.9	80.75	77.2	8.33
1.2	4	199.1	35.72	71	7.56
1.2	6	133.2	19.59	67.8	6.43
1.5	1	31.4	0.8	14	0
1.5	2	21.7	0.46	14	0
1.5	4	14.9	0.3	14	0
1.5	6	14	0	14	0
2	1	21	0.45	11	0
2	2	14.9	0.3	11	0
2	4	11.3	0.46	11	0
2	6	11	0	11	0
3	1	63	14.2	26.8	5
3	2	43.5	8.8	21.9	3.81
3	4	22.1	2.74	18.7	1.68
3	6	19.3	1.68	18.4	1.36
6	1	1,499.3	1,879.2	445.3	376.56
6	2	691.7	744.38	262.4	183.25
6	4	200.1	90.35	99.2	38.2
6	6	99.2	31.99	56.9	14.94
10	1	3,973.9	5,692.81	1,021.7	1,049.87
10	2	1,413.8	1,399.17	563.4	507.32
10	4	396.6	195.55	208.1	82.09
10	6	188.3	64.18	112	33.12

a full exemplary range of ℓ_p spaces with $p \in \{1.1, 1.2, 1.5, 2, 3, 6, 10\}$ and the power type of the duality mapping chosen as before.

The average iteration counts and standard deviations are given in Table 2.

First of all, we notice that with $p \in \{1.2, 1.5, 2\}$ there is not much change in the average iteration count n over different number of search directions N. This holds on for $p = 1.5$ within output precision of 10^{-7} and for $p = 2$ within full floating point numerical precision. Hence, we see that when $\mathcal{X} = \ell_2(\mathbb{R}^m)$ is a Hilbert-space, the orthogonality is maintained between all search directions from $\mathbf{v}^*_{0,N_0}$ up to $\mathbf{v}^*_{n,N_n}$. This is the behavior expected from CG methods and we elucidate this further in the next section. It is maintained to some extent when p is close to 2.

Second, for $p \in \{1.5, 2\}$ we observe that for large number of search directions N the average number of iteration steps n becomes equivalent for both of the

search spaces $\mathcal{U}_n^{\text{trunc}}$ and $\mathcal{V}_n^{\text{trunc}}$. This indicates that using more than one search direction, the central idea of [21], is indeed substantial.

And last, iteration counts and deviations become very large for $p \to 1$ and $p \to \infty$, i. e. when the ℓ_p spaces are no longer smooth.

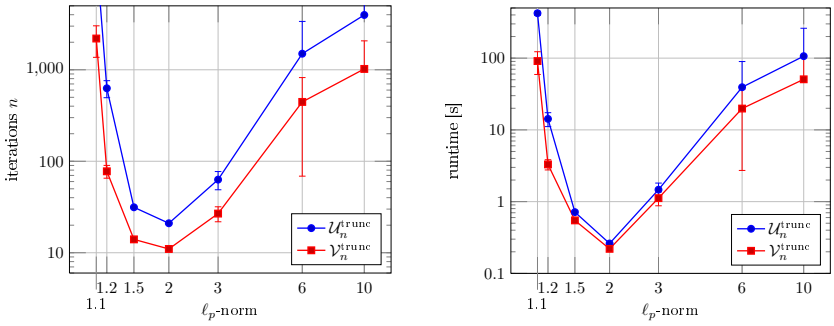

(a) Iterations (b) Runtimes

Figure 2: Iterations and runtimes solving with SESOP and two different search spaces for the minimum-norm solution of $\mathbf{Ax} = \mathbf{y}$ with uniformly random $\mathbf{A} \in \mathbb{R}^{1000\times5000}$ and a single search direction.

In Figure 2 we depict for a single search direction both iteration counts and the total runtime for solving for the minimum-norm solution up to a relative residual threshold of 10^{-4} or up to 20.000 iteration steps. Here, we want to compare the method's performance with either search space directly.

We notice that iteration counts for the orthogonalized search space $\mathcal{V}_n^{\text{trunc}}$ are at least a factor of 2-3 below the ones for the search space $\mathcal{U}_n^{\text{trunc}}$. This holds over all values of p. This reduced number of iterations required for the same residual threshold is the reason, why despite the additional computational effort for the orthogonalization, also the runtimes show the same trend between the compared search spaces up to a similar factor.

Finally, we look at runtimes for multiple search directions in Figure 3.

We see that for more than one search direction N the additional cost for the orthogonalization procedure of the search directions at some point starts to outweigh the gain obtained by a reduced number of total iterations. There are two trends behind this: First, SESOP with an unorthogonalized search space also becomes more effective when using multiple search directions, c. f. Table 2. Second, the overhead of orthogonalization becomes more costly when N increases. These two trends work especially in favor of SESOP with $\mathcal{U}_n^{\text{trunc}}$ for p around 2 where the required iteration steps are comparable for $N \in \{4, 6\}$ and explains why the unorthogonalized procedure is then faster in the total runtimes lacking the additional orthogonalization. This holds for $p \in \{1.5, 2, 3\}$.

Looking at the trend in runtimes for $\mathcal{V}_n^{\text{trunc}}$ for increasing number of search directions N_n, we see little to no gain. Runtimes for $p \leq 2$ generally increase. This trend turns around for towards $p > 2$ where the runtimes decrease when using more search directions.

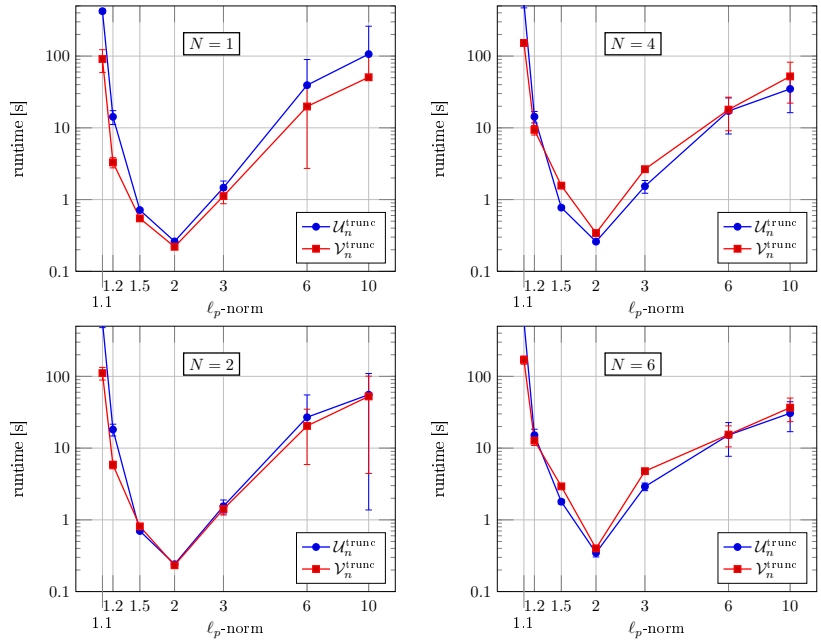

Figure 3: Runtimes solving with SESOP and two different search spaces for the minimum-norm solution of $\mathbf{Ax} = \mathbf{y}$ with uniformly random $\mathbf{A} \in \mathbb{R}^{1000\times 5000}$ and $N \in \{1, 2, 4, 6\}$.

4.1.4 Connection to CG in Hilbert Space

We briefly want to review the ℓ_2-case, where the Banach space is a Hilbert space. There, equation (31) is identical to the update used by the Polak-Ribière [16] type of CG method on the normal equation,

$$\mathbf{Bx} = \mathbf{z} \quad \text{with } \mathbf{B} = \mathbf{A}^*\mathbf{A}, \mathbf{z} = \mathbf{A}^*\mathbf{y}.$$

Hence, we expect that the conjugacy property will hold in our case, too. Note that it takes the following form with $\mathbf{p}_i = \mathbf{Ax}_i - \mathbf{y}$,

$$\langle \mathbf{A}^*\mathbf{p}_i, \mathbf{A}^*\mathbf{p}_j \rangle = 0 \quad \forall i \neq j. \tag{44}$$

Furthermore, we have seen in Table 2 that there is no difference between a single search direction and multiple search directions as the current gradient direction is made orthogonal to all previous search directions simultaneously.

In Figure 4 we give the Bregman distance to the in our case known true solution $\widehat{\mathbf{x}}$ per iteration step for three different ℓ_p norms, one of them being the ℓ_2 norm. There, we only look at the orthogonalized search space $\mathcal{V}_n^{\text{trunc}}$ using metric projections.

We clearly notice that in the ℓ_2 case more than a single search direction does not change anything about the minimization. It proceeds (up to numerical precision) in exactly the same manner as if there were only a single search direction spanning the search space. This is precisely what would be expected of a CG method due to the inherent conjugacy property. This is not valid for the other two cases with $p = 1.2$ and $p = 6$ where more search directions clearly further aid convergence, especially for $p \gg 2$.

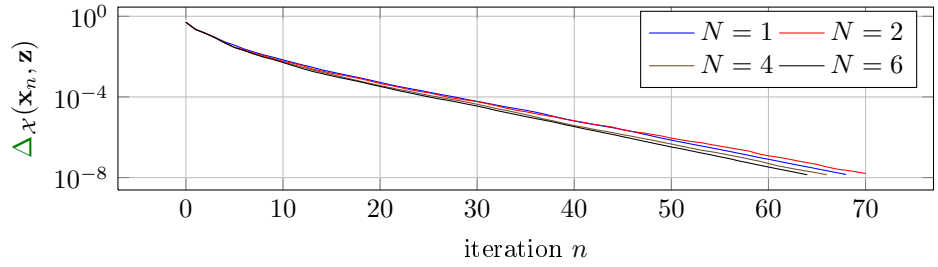

(a) $p = 1.2$

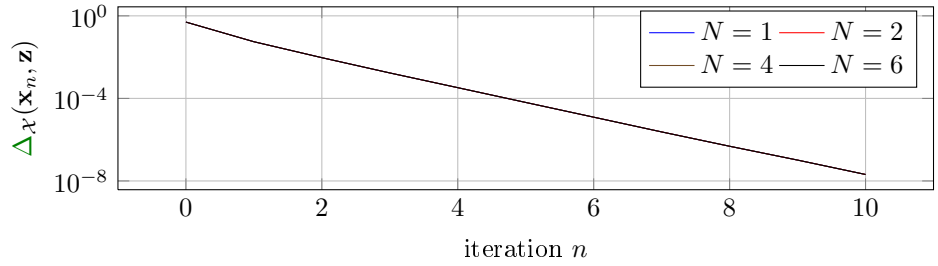

(b) $p = 2$

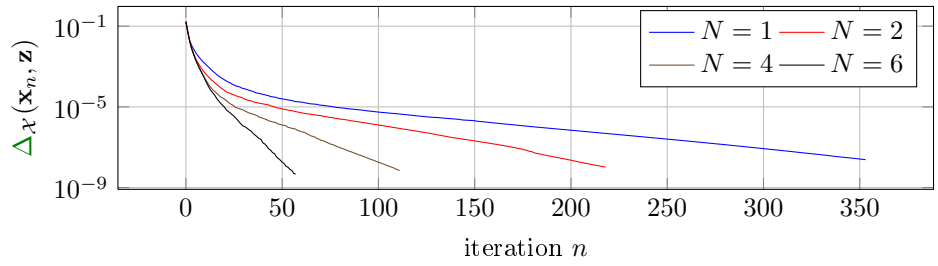

(c) $p = 6$

Figure 4: Bregman distance $\Delta_{\mathcal{X}}(\mathbf{x}_n, \mathbf{z})$ over the iteration step n for SESOP with search space $\mathcal{V}_n^{\text{trunc}}$ for increasing number of search directions N, for various ℓ_p norms with $m = 5000$, and for a specific random number generator seed.

Hence, we also see in this numerical example that the subspace methods (with orthogonalization) in fact are an extension of the CG methods to general Banach spaces.

4.2 Computerized Tomography

The computerized tomography problem in 2D aims at reconstruction of the inside of an object from projections. Measurements are for example obtained by passing radiation through a body, whereby their intensity is diminished, proportional to the passed length and density of the body $f(x) : [0, 1]^2 \to \mathbb{R}_0^+$. This decrease is measured over a a angles and s shifts of radiation source and detectors. This measurement matrix is usually called *sinogram*. We follow [9, Sect. 7.7] in the brief introduction to the discretized setting.

The measurement rays are parametrized as

$$t^i(\tau) = t^{i,0} + \tau d^i, \tag{45}$$

where d is the directional vector of the ith ray with $i = 1, \dots, s \cdot a$.

Then, the dampening of the ith ray can be written as the *Radon transform*,

$$b_i = \int_{-\infty}^{\infty} f\left(t^i(\tau)\right) d\tau, \quad i = 1, \dots, s \cdot a, \tag{46}$$

i. e. we integrate $d\tau$ along the line $t^i(\tau)$.

The problem can be discretized with a pixel basis,

$$\chi_{kl}(x) = \begin{cases} 1, & x \in [h \cdot (k-1), h \cdot k] \times [h \cdot (l-1), h \cdot l] \\ 0, & \text{else} \end{cases}, \quad \forall k, l = 1, \dots, n, \tag{47}$$

where we constrain the absorption coefficient $f(x) = \sum_{kl} f_{kl} \chi_{kl}(x)$ to be piecewise constant with $h = \frac{1}{n}$.

For the discretization of (46) we simply need to count the length $\Delta L_{kl}^{(i)}$ of each ray t^i in each pixel χ_{kl} of the basis (47) and obtain

$$b_i = \sum_{k,l=1}^{n} f_{kl} \Delta L_{kl}^{(i)} \quad \text{for } i = 1, \dots, s \cdot a. \tag{48}$$

If we vectorize the matrix object f_{kl} to become the vector x_j with $j = (l-1)n + k$, then we obtain

$$b_i = \sum_j a_{ij} x_j, \quad i = 1, \dots, s \cdot a. \tag{49}$$

Note that the matrix A is sparse what we use in the implementation.

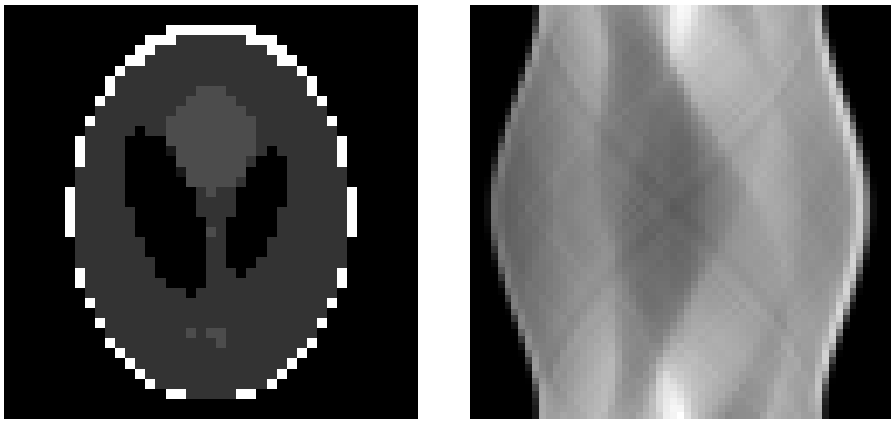

(a) Phantom (b) Sinogram

Figure 5: Shepp-Logan phantom and its analytically obtained sinogram.

We look at the standard Shepp-Logan phantom, see Figure 5(a), where we obtain the measurements, see Figure 5(b), by using the known analytical Radon transform of ellipses. We look at only a single exemplary discretization with number of pixels $n = 41$, number of shifts $s = 61$, and number of discrete angles $a = 60$. We deliberately choose a coarser image resolution together with a higher number of measurements to allow for a high-quality reconstruction and clearly discernable artifacts if there are any.

Note that we additionally project the solution onto the range of the matrix A such that the true solution can be computed by the method. We return to this point in Section 4.2.2.

Here, we again want to compare SESOP using the truncated search space $\mathcal{U}_n^{\text{trunc}}$ and also the orthogonalized search space $\mathcal{V}_n^{\text{trunc}}$. They have the following dimensions: $|\mathcal{U}_n^{\text{trunc}}| = 2$ and $|\mathcal{V}_n^{\text{trunc}}| = 1$. We stop the methods either after 500 iterations or when the absolute residual has decreased below 10^{-2}. We use ℓ_p-spaces with $p \in \{1.1, 1.2, 1.5, 2\}$ for $\mathcal{X}$ and ℓ_2 for $\mathcal{Y}$. For the line-search problem we use at most 20 iterations. The power type of the duality mappings is always set to 2.

4.2.1 Exact data

First of all, we look at the problem without noise and only in the case of $\mathcal{X} = \ell_2$, i. e. the Hilbert space setting.

In Figure 6 we give the recovered solution and the error, residual, and Bregman distance histories for SESOP with the truncated search space $\mathcal{U}_n^{\text{trunc}}$ and two search directions.

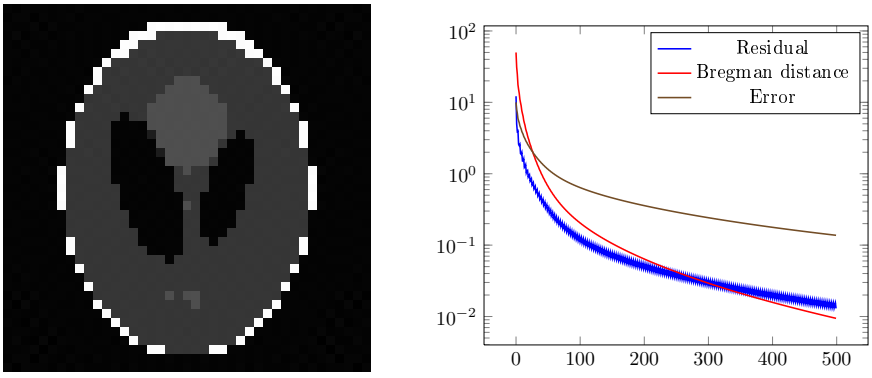

(a) solution (b) iteration history

Figure 6: Recovered solution and residual, error, and Bregman distance histories using SESOP with $\left|\mathcal{U}_n^{\text{trunc}}\right| = 2$ directions.

We see that SESOP converges but stops at $n = 500$ iterations where the absolute residual is still slightly larger than 10^{-2}. The overall runtime is 2.4 seconds. Note that the convergence is monotone only in the Bregman distance and not in the residual, c. f. Thm 2.

Next, we look at SESOP using the orthogonalized search space $\mathcal{V}_n^{\text{trunc}}$ with a single search direction[3]. We obtain the results as given in figure 7.

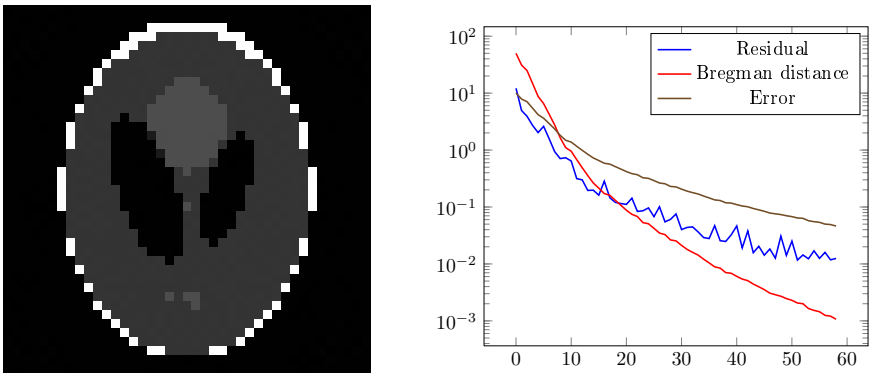

(a) solution (b) iteration history

Figure 7: Recovered solution and residual, error, and Bregman distance histories using orthogonalized SESOP with $\left|\mathcal{V}_n^{\text{trunc}}\right| = 1$ direction.

Here, we notice that the iteration stops at roughly $n = 60$ when reaching an absolute residual of 10^{-2}. Naturally, this translates to a faster runtime compared to the unorthogonalized search space of only 0.36 seconds.

In the former case the reconstructed image retains some very slight artefacts, in the latter case no artefacts are visible.

4.2.2 Noisy data

It is called *inverse crime*, see [9, Sect. 7.2], if the same discretization has been used for both the operator and the right-hand side. Results may look suspiciously good in this case, i. e. no mismatch between (real) data and model is revealed. We definitely committed this crime by projecting our measurements onto the range of the matrix A. Now, we additionally disturb the right-hand side $\mathbf{y}$, obtained from projecting the Shepp-Logan phantom $\widehat{\mathbf{x}}$ with the matrix $\mathbf{y} = \mathbf{A}\widehat{\mathbf{x}}$, with noise of a known level δ to become $\mathbf{y}^\delta$. This ensures that $\left\|\mathbf{y} - \mathbf{y}^\delta\right\|$ is known to us.

To this end, we construct a random vector $\mathbf{n} \in [-1, 1]^m$ and set $\widetilde{\mathbf{y}} = \mathbf{y} + \delta\frac{\|\mathbf{y}\|}{\|\mathbf{n}\|}\mathbf{n}$ with the projected right-hand side $\mathbf{y} = \mathbf{A}\widehat{\mathbf{x}}$ using the true solution $\widehat{\mathbf{x}}$, which is the Shepp-Logan phantom. We use a noise level of $\delta = 0.01$. We stop the iteration via a relative residual threshold of 0.01 with a discrepancy parameter of 3, i. e. the relative residual threshold becomes 0.03.

As a perfect reconstruction is in the noisy case no longer possible, we increase the dimensions of the measurement matrix to number of offsets of $s = 81$ and number of angles to $a = 80$. Also we increase the number of pixels $p = 81$ of the reconstructed image. We still remain in the Hilbert space setting $\mathcal{X} = \ell_2$ for the moment. In the next section we look at other ℓ_p spaces.

[3] As mentioned before, in an ℓ_2 space more search directions do not improve performance when using an orthogonalized search space.

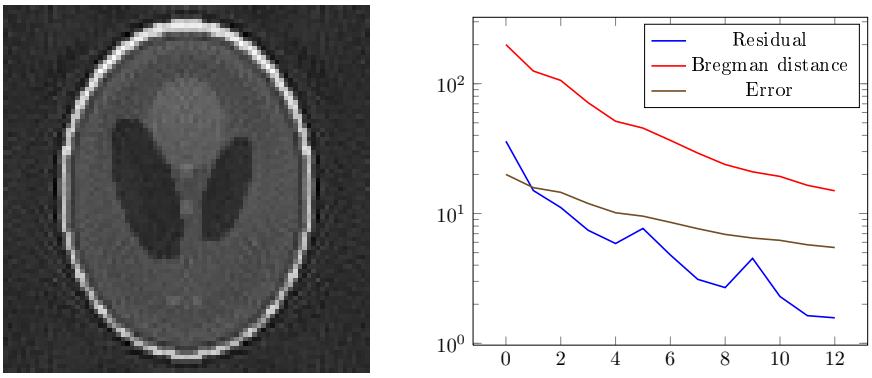

(a) solution (b) iteration history

Figure 8: Recovered solution and residual, error, and Bregman distance histories using orthogonalized SESOP with $\left|\mathcal{V}_n^{\text{trunc}}\right| = 1$ direction in the presence of noise of level 0.01 and using a stopping criterion with $\tau = 3$.

In Figure 8 both the reconstructed image and iteration history with residual, Bregman distance, and error with respect to the true solution is shown. We obtain good results with respect to the noise level employed. The runtime is 0.34 seconds.

We conclude that we do not commit any inverse crime and that the proposed method is indeed working and implemented sensibly.

4.2.3 Small ℓ_p Norms

Finally, we also look at other norms than ℓ_2, namely small ℓ_p norms with $p \in \{1.1, 1.2, 1.5\}$. Here, to allow for a high-quality reconstruction we again use $s = 61$, $a = 60$ and $p = 41$.

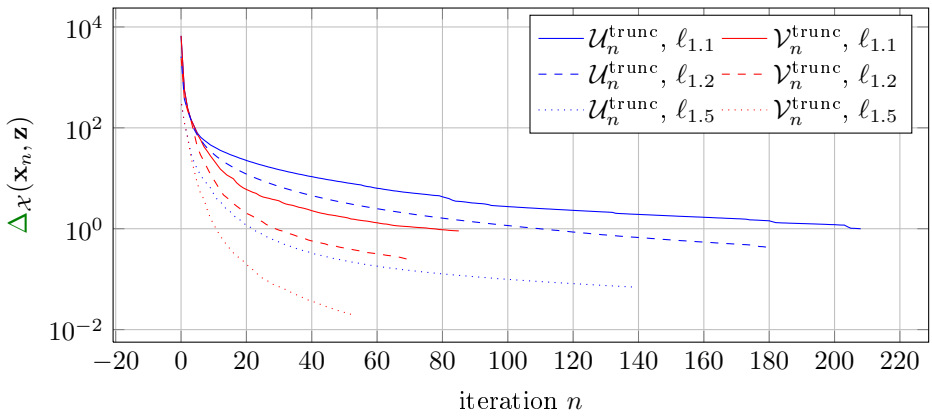

Figure 9: Bregman distance $\Delta_{\mathcal{X}}(\mathbf{x}_n, \mathbf{z})$ over the iteration step n for SESOP with unorthogonalized search space $\mathcal{U}_n^{\text{trunc}}$ and orthogonalized search space $\mathcal{V}_n^{\text{trunc}}$ in comparison on the norms $\ell_{1.1}$, $\ell_{1.2}$, and $\ell_{1.5}$.

In Figure 9 we compare the decrease in Bregman distance for unorthogonalized and orthogonalized search spaces for various small ℓ_p norms. We see that over all tested small ℓ_p norms the orthogonalized search space $\mathcal{V}_n^{\text{trunc}}$ requires about three times fewer iterations to reach the stopping criterion. This makes it also faster in the overall runtimes, e. g. for $p = 1.1$ we obtain 3.15 seconds without orthogonalization and 2.6 seconds when employing $\mathcal{V}_n^{\text{trunc}}$. Note that because of different ℓ_p used to measure the residual stopping criterion, the method also terminates at different Bregman distances.

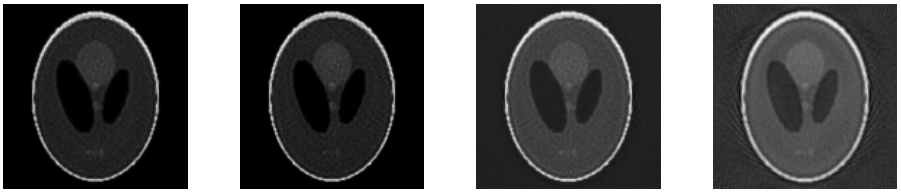

(a) $\ell_{1.1}$ (b) $\ell_{1.2}$ (c) $\ell_{1.5}$ (d) ℓ_2

Figure 10: Reconstructed Shepp-Logan phantom in the presence of noise of level $\delta = 0.01$ and using small ℓ_p norms with $p \in \{1.1, 1.2, 1.5, 2\}$.

Last of all, we show in Figure 10 larger reconstructed images for $p = 127$ pixels with a noise of $\delta = 0.01$ while using same number of angles a and shifts s than before. This is to elucidate the effect of the different ℓ_p norms in the presence of noise.

We observe that for smaller ℓ_p both contrast and also noise of the image is enhanced. Especially, artifacts surrounding the reconstructed phantom in the ℓ_2 case are absent for $\ell_{1.1}$. On the other hand, noisy speckles are more present in the reconstructed image using the $\ell_{1.1}$ norm.

5 Conclusions

Based on the previous work of [18], we have proposed a semi-orthogonalized set of search directions in Banach spaces using metric projections. Using search spaces consisting only of the last Landweber descent directions modified by this orthogonalization, we have shown that the Sequential Subspace Optimization (SESOP) method converges weakly and that for the Hilbert space case the procedure coincides with the Polak-Ribière Conjugate Gradient (CG) method applied to the normal equation. Hence, the subspace methods are a natural extension to general Banach spaces of CG methods, that are known to work very efficiently in Hilbert spaces.

Numerical experiments have shown fast convergence for both, an inverse toy problem consisting of a uniformly distributed random matrix and right-hand side on various ℓ_p spaces as well as for the inverse problem of 2D computerized tomography using the discretized backprojection as operator and a Radon transformed Shepp-Logan phantom as right-hand side. In every case the orthogonalized truncated search space clearly outperforms the truncated search space used in [21], both in required iterations and runtime.

Note that although we have investigated the case of a noisy right-hand side in the experiments, convergence for this case has not been done yet.

As an outlook, we would like to remind that there is a whole zoo of CG variants, see [1, p. 98]. It would be very insightful to find more connections between a specific variant and a choice of (semi-orthogonalized) search spaces for the SESOP method. Proving the regularization property of the method is another object of future research.

Acknowledgements

The authors acknowledge funding by the German ministry of Education and Research (BMBF) under the project HYPERMATH (05M2013).

References

[1] N. Andrei, *Another nonlinear conjugate gradient algorithm for unconstrained optimization*, Optimization Methods and Software, 24 (2009), pp. 89–104, doi:10.1080/10556780802393326.

[2] Y. Censor and A. Zenios, *Parallel Optimization - Theory, Algorithms, and Applications*, Oxford University Press, New York, Oxford, 1st ed., 1997.

[3] C. Chidume, *Geometric Properties of Banach Spaces and Nonlinear Iterations*, vol. 1965 of Lecture Notes in Mathematics, Springer London, London, 2009, doi:10.1007/978-1-84882-190-3.

[4] I. Cioranescu, *Geometry of Banach Spaces, Duality Mappings and Nonlinear Problems*, Kluwer Academic Publishers, Dordrecht, The Netherlands, 1. ed., 1990.

[5] J. A. Clarkson, *Uniformly Convex Spaces*, Transactions of the American Mathematical Society, 40 (1936), p. 396, doi:10.2307/1989630.

[6] M. M. Day, *Uniform Convexity in Factor and Conjugate Spaces*, The Annals of Mathematics, 45 (1944), p. 375, doi:10.2307/1969275.

[7] N. Dunford and J. T. Schwartz, *Linear Operators - Part I - General Theory*, Interscience Publishers, Inc., New York, 1st ed., 1957.

[8] G. Guennebaud, B. Jacob, and Others, *Eigen v3*, 2010, http://eigen.tuxfamily.org.

[9] P. C. Hansen, *Discrete Inverse Problems - Insight and Algorithms*, SIAM, 1st ed., 2010.

[10] J. Lindenstrauss, *On the modulus of smoothness and divergent series in Banach spaces.*, The Michigan Mathematical Journal, (1963).

[11] J. Lindenstrauss and L. Tzafriri, *Classical Banach Spaces II - Function Spaces*, Springer-Verlag Berlin Heidelberg New York, 1. ed., 1979.

[12] I. Loris, *On the performance of algorithms for the minimization of l1 - penalized functionals*, Inverse Problems, 25 (2009), p. 035008, doi:10.1088/0266-5611/25/3/035008.

[13] J. Muscat, *Functional Analysis*, Springer International Publishing, Cham, 1st ed., 2014, doi:10.1007/978-3-319-06728-5.

[14] G. Narkiss and M. Zibulevsky, *Sequential Subspace Optimization Method for Large-Scale Unconstrained Problems.* 2005.

[15] J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer-Verlag, New York, 1. ed., 1999.

[16] E. Polak and G. Ribiere, *Note sur la convergence de méthodes de directions conjuguées*, ESAIM: Mathematical Modelling and Numerical Analysis - Modélisation Mathématique et Analyse Numérique, 3 (1969), pp. 35–43.

[17] M. Schechter, *Principles of functional analysis*, Academic Press, Inc., New York, 1st ed., 1971.

[18] F. Schöpfer, *Iterative Regularization Methods for the Solution of the Split Feasibility Problem in Banach Spaces*, doctoral thesis, Universität des Saarlandes, 2007.

[19] F. Schöpfer, A. K. Louis, and T. Schuster, *Nonlinear iterative methods for linear ill-posed problems in Banach spaces*, Inverse Problems, 22 (2006), pp. 311–329, doi:10.1088/0266-5611/22/1/017.

[20] F. Schöpfer and T. Schuster, *Fast regularizing sequential subspace optimization in Banach spaces*, Inverse Problems, 25 (2009), pp. 1–22, doi:10.1088/0266-5611/25/1/015013.

[21] F. Schöpfer, T. Schuster, and A. K. Louis, *Metric and Bregman projections onto affine subspaces and their computation via sequential subspace optimization methods*, Journal of Inverse and Ill-posed Problems, 16 (2008), pp. 1–29, doi:10.1515/JIIP.2008.026.

[22] Z.-b. Xu and G. F. Roach, *Characteristic inequalities of uniformly convex and uniformly smooth Banach spaces*, Journal of Mathematical Analysis and Applications, 157 (1991), pp. 189–210, doi:10.1016/0022-247X(91)90144-O.

[23] Y.-X. Yuan and J. Stoer, *A subspace study on conjugate gradient algorithms*, ZAMM - Journal of Applied Mathematics and Mechanics, 75 (1995), pp. 69–77.

[24] E. Zeidler, *Nonlinear functional analysis and its applications. 1. Fixed-point theorems*, Springer-Verlag Berlin Heidelberg New York, Berlin Heidelberg New York, 1st ed., 1986.